

Contributions to measures of segregation or polarization and to nonasymptotic inference in linear models

Lucas Girard

► To cite this version:

Lucas Girard. Contributions to measures of segregation or polarization and to nonasymptotic inference in linear models. Economics and Finance. Institut Polytechnique de Paris, 2021. English. NNT: 2021IPPAG007. tel-03555582

HAL Id: tel-03555582 https://theses.hal.science/tel-03555582

Submitted on 3 Feb 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.





Contributions to measures of segregation or polarization and to nonasymptotic inference in linear models

Thèse de doctorat de l'Institut Polytechnique de Paris préparée à l'École Nationale de la Statistique et de l'Administration Économique (ENSAE)

> École doctorale n°626 de l'Institut Polytechnique de Paris (ED IP Paris) Spécialité de doctorat : Sciences économiques

> > Thèse présentée et soutenue à Palaiseau, le 3 décembre 2021, par

LUCAS GIRARD

Composition du Jury :

Roland Rathelot Professeur, ENSAE-CREST	Président
Maximilian Kasy Professeur associé, Oxford University	Rapporteur
Daniel Wilhelm Professeur associé, University College London	Rapporteur
Russell Davidson Professeur, McGill University	Examinateur
Xavier D'Haultfœuille Professeur, ENSAE-CREST	Directeur de thèse

Thèse de doctorat

Remerciements – Acknowledgments

Ce manuscrit n'existerait pas sans la présence, le soutien et la confiance des nombreuses personnes qui m'ont permis de débuter puis d'achever cette thèse. Ces quelques mots visent à les remercier.

First of all, I thank Russell Davidson, Maximilian Kasy, Roland Rathelot, and Daniel Wilhelm for accepting to be members of the jury and to assess my work. I thank, in particular, Maximilian Kasy and Daniel Wilhelm for being referees of my thesis and for their comments.

Je remercie Xavier D'Haultfœuille, mon directeur de thèse, pour sa disponibilité, sa confiance et sa bienveillance, notamment durant certaines périodes de doutes et d'incertitudes. Je me souviens notamment de nos échanges à des heures tardives, fin avril 2017, pour finaliser mon dossier de candidature à la bourse doctorale du GENES, puis de son écoute et de ses conseils à l'été 2017, où, alors en stage de recherche au CREST, j'hésitais à me lancer dans la thèse. Je n'aurais pu réaliser celle-ci sans son aide, son soutien et ses nombreuses idées, toujours pertinentes.

Je remercie également Xavier pour son enseignement en économétrie et tout ce qu'il m'a permis de comprendre à cette discipline un peu mystérieuse par certains aspects, à l'intersection des statistiques et de l'économie (ou plus largement, des sciences sociales). Tout d'abord, en tant qu'étudiant à l'ENSAE en deuxième et troisième années pour une initiation. Puis, pour une compréhension progressivement plus aboutie, bien qu'encore fort partielle, en tant que doctorant et chargé de Travaux Dirigés (TD) pour ses cours d'Économétrie 1 et 2. J'espère avoir pu moi-même transmettre aux étudiants les points principaux de ses cours ambitieux et formateurs.

Je remercie mes co-auteurs, mon directeur de thèse Xavier et Roland Rathelot avec lesquels j'ai beaucoup appris, dans le fond et le contenu mais également dans les a-côtés (forme, choix de présentation, choix de soumission, etc.) d'un projet de recherche, lesquels sont loin d'être négligeables en vue d'une publication. Je remercie également Roland pour son suivi et ses conseils durant l'ensemble de ma thèse.

Je remercie mes co-auteurs Alexis Derumigny et Yannick Guyonvarch. Tous trois encore doctorants au CREST dans l'ancien bureau d'Alexis au pôle Finance ou désormais en ligne, chacun de nous sur l'Overleaf partagé, j'apprécie toujours autant nos réunions de travail et d'avoir la chance d'échanger sur nos projets ou sur d'intéressantes digressions, comme les mystères toujours renouvelés des modèles linéaires.

Bien que ce projet ne fasse finalement pas partie de ce manuscrit, je remercie également Bruno Crépon, Anett John et Esther Mbih, avec lesquels j'ai pu m'initier à la conception d'un questionnaire et à la réalisation d'un projet avec Pôle emploi.

Je remercie aussi Bruno ainsi qu'Alessandro Riboni pour m'avoir suivi et conseillé lors de mes comités de thèse et, plus largement, l'ensemble des professeurs du CREST pour leur disponibilité et leur aide. Je remercie Laurent Davezies pour m'avoir fait confiance et engagé comme chargé d'enseignement à la direction des études de l'ENSAE. Ce poste m'a permis, moyennant quelques heures de TD et quiz d'économétrie supplémentaires, de finir ma thèse dans les meilleures conditions. Si cela ne transparaît pas directement dans ce manuscrit, l'enseignement a occupé une place importante dans mon doctorat, en parallèle de la recherche. C'est une activité que j'apprécie, dont la portée pourrait être décrite comme plus immédiate, plus directe par rapport à la recherche, et qui en cela a un aspect réconfortant par rapport aux péripéties d'un projet de recherche. Je remercie ainsi Laurent et la direction de l'ENSAE, le directeur de l'école Pierre Biscourp, le successeur de Laurent à la direction des études, Frédéric Loss, et la directrice-adjointe des études Rosalinda Solotareff, de m'avoir donné l'opportunité de poursuivre et d'approfondir cette activité en parallèle de la partie recherche de mon doctorat. Donner des TD est probablement la meilleure façon de se rendre compte du peu qu'on avait compris étudiant, et de comprendre un peu plus chaque année !

Je remercie Laurent aussi pour son enseignement. En troisième année de l'ENSAE, bien que je ne suivisse pas ce cours, j'assistai à une séance du "Séminaire de sondages", dans une petite salle de l'ancien bâtiment accolé à la tour de l'INSEE à Malakoff, où Laurent discuta de la prise en compte des poids de sondages en économétrie. Ce cours et la note qui lui est liée, écrite avec Xavier, reste pour moi un modèle d'économétrie théorique en apportant une réponse fondée, complète et précise à une question méthodologique souvent importante en pratique. Dans cette lignée, je remercie Victor-Emmanuel Brunel, professeur de statistiques à l'ENSAE. Durant mon doctorat, j'ai pu suivre ses deux cours de troisième année de l'ENSAE. Son enseignement, à l'ancienne si j'ose dire, intégralement au tableau et à la craie tout en ajoutant à l'oral la motivation et les remarques pertinentes reste pour moi un exemple pédagogique.

Je remercie le CREST et ses directeurs, Francis Kramarz puis Arnak Dalalyan, pour m'avoir permis d'engager cette thèse en m'octroyant une bourse doctorale. Je remercie l'ENSAE et l'ensemble de la direction des études pour m'avoir permis de la poursuivre, un étage en dessous mais toujours dans un environnement agréable. Agréable et également fonctionnel grâce à l'équipe informatique que je remercie pour l'aide tant logicielle que matérielle, et grâce à l'équipe administrative que je remercie également, notamment Pascale Deniau et Edith Verger pour leur aide dans l'organisation du séminaire de microéconométrie du CREST durant ma deuxième année de thèse.

Je remercie mes co-bureaux du 4097, Germain Gauthier, Bérengère Patault et Anasuya Raj, puis également Elie Gerschel, pour leurs conseils, leurs discussions, et pour avoir rendu mon quotidien de doctorant si agréable et plaisant, tout en restant (parfois) studieux.

Depuis mon déménagement au troisième étage, j'ai eu la chance de poursuivre de telles discussions et moments de pause avec les coordinateurs d'enseignement et les autres membres de la direction des études, notamment Rémi Avignon, Morgane Cure, Jules Depersin, Fabien Perez et Jérôme Trinh. Je pense en particulier à nos discussions avec Fabien et à nos essais, toujours renouvelés (et parfois fructueux), de faire comprendre au mieux les cours d'économétrie aux élèves de l'ENSAE. Je les remercie pour leur aide et toutes ces petites choses pourtant indispensables pour apprécier la thèse au quotidien : une pause déjeuner-natation à la piscine d'Orsay avec Morgane, un retour en vélo par la vallée de la Bièvre avec Jules et Rémi, un karaoké-raclette chez Jérôme.

Pour ces mêmes raisons, je remercie grandement l'ensemble des doctorants que j'ai eu la chance de rencontrer durant ma thèse pour leur soutien et pour avoir rendu le CREST un endroit où j'avais plaisir à venir et à travailler tous les jours. Je crains que ma mémoire ne me fasse défaut mais je pense notamment (dans un ordre plus ou moins chronologique) à Arthur Cazaubiel, Christophe Gaillac, Jérémy L'Hour, Ivan Ouss, Julie Pernaudet, Benjamin Walter, Ao Wang, aux membres de l'équipe du ManCRESTer United, Reda Aboutajdine, Antoine Bertheau, Aurélien Bigo, Wael Bousselmi, Gwen-Jiro Clochard, Thomas Delemotte, Etienne Guigue, Jérémy Hervelin, Corentin Lechaux, Clémence Lenoir, Hugo Molina, Julien Monardo, Martin Mugnier, Antoine Valtat et Tang Yuanzhe, aux rivaux (appréciés) du CRESTiano, Badr-Eddine Chérief-Abdellatif et Geoffrey Chinot, ainsi qu'à Léa Bou Sleiman, Pauline Carry, Héloïse Cloléry, Pierre-Edouard Collignon, Antoine Ferey, Alice Lapeyre, Alexis Larousse, Elio Nimier-David, Louis Pape, Félix Pasquier, Elia Pérennès, Inès Picard, et Emilie Sartre.

Puisque la thèse marque officiellement la fin des études, je souhaite remercier les amis qui m'ont accompagné à chaque marche de ce long escalier : Antoine et Martin depuis le tout début ou presque ainsi que Javed et Olivia, nos tennis hebdomadaires du samedi, que ce soit sous la neige, la pluie, voire parfois le soleil, auront toujours été l'occasion de se ressourcer entre amis, notamment dans certains moments de doute ; Corentin au lycée ; Paul M durant ma prépa ; Xavier, Quentin, Luc, Alexis, Nino et Guillaume à HEC ; Marie-Aude, Pauline, Ismaïl, Aurélien et Ashley lors de la licence de maths à Orsay en parallèle ; et Paul V et Satya durant mes années à l'ENSAE. Je remercie également mes amis du badminton, notamment Vincent, Robert et Adrien, et ceux du foot.

Je tiens aussi à remercier l'ensemble de mes professeurs, entraîneurs et éducateurs depuis mon enfance. Je ne saurais tous les nommer ici mais je garderai toujours un souvenir ému et reconnaissant de plusieurs d'entre eux.

Je remercie Laure pour m'avoir accompagné (et supporté) durant ces quatre années de thèse. Je n'y serais jamais parvenu sans ton aide et tous les bons moments passés ensemble.

Enfin, je remercie mes parents, mes sœurs, mes grands-pères, mes oncles, le reste de ma famille recomposée et Laure pour leur soutien et amour inconditionnels qui rendent ma vie merveilleuse en tout et malgré tout et en font sa valeur. Je remercie mes parents qui m'ont toujours soutenu dans mes choix d'études malgré mes hésitations. Je pense aussi à mes grand-mères à qui je dédie cette thèse.

L.G., le 29 novembre 2021

Contents

Su	imm	ary in English	5				
Re	Résumé substantiel en français						
In	Introduction in English 9						
In	trod	uction en français	21				
1	Mea	asuring segregation on small units	37				
	1.1	Introduction	38				
	1.2	Set-up, estimation, and inference	39				
		1.2.1 The setting and the parameters of interest	39				
		1.2.2 Nonparametric approach	40				
		1.2.3 Parametric approach	44				
		1.2.4 Correction of the naive index	45				
		1.2.5 Test of the binomial assumption	46				
		1.2.6 Conditional segregation indices	46				
	1.3	The segregsmall command	48				
		1.3.1 Syntax	48				
		1.3.2 Description and main options	48				
		1.3.3 Additional options	49				
		1.3.4 Saved results	50				
		1.3.5 Execution time	53				
	1.4	Example	54				
	1.5	Conclusion	57				
	App	pendix 1.A Magnitude of the small-unit bias	59				
	App	pendix 1.B Plug-in estimators of naive indices	60				
	App	pendix 1.C Indices in the parametric approach	60				
	App	pendix 1.D Supplements to the example	61				
2	Mea	asures of residential segregation in France	65				
	2.1	Introduction	66				
	2.2	Data and dimensions	72				
		2.2.1 Labor Force Survey and units/neighborhoods	72				
		2.2.2 Dimensions of segregation and individual covariates	73				
	2.3	Methodology and robustness	77				
		2.3.1 Parameters of interest, estimation, and inference	77				
		2.3.2 Specification choices and robustness	79				
	2.4	Results	82				

		2.4.1	Unconditional analysis	. 84
	٨	2.4.2	Conditional analysis	. 89
	App	endix 2.	A Additional figures	. 93
3	Ide	ntificati	on and estimation of speech polarization	103
	3.1	Introdu	\mathbf{r}	. 104
	3.2	The fra	umework	. 108
		3.2.1	Statistical model	. 109
		3.2.2	Parameter of interest	. 110
	3.3	Main t	heoretical results	. 112
		3.3.1	Bounds under minimal conditions	. 112
		3.3.2	Point-identification and estimation through extrapolation	. 114
	3.4	Extens	$\overline{\mathrm{ons}}$. 115
		3.4.1	Including covariates	. 115
	2 5	3.4.2	Test of the binomial assumption	. 117
	3.5	Applica	ation to speech polarization in the U.S. Congress $(1873-2016)$. 118
		3.3.1	Distinguishing processing operations and small unit bias	. 118
		3.5.2	Testing our approach	. 119
	3.6	Conclu	sion	. 122
	App	endix 3	A Proofs of main theorems	120
	- PP	3.A.1	Proof of Theorem 3.1 (definition of the index)	. 129
		3.A.2	Proof of Theorem 3.2 (partial identification)	. 130
	App	endix 3.	B Supplements to the methodology	. 136
		3.B.1	Connection with the Coworker index of segregation	. 136
		3.B.2	Restriction to $K > 0$ without loss of generality	. 138
		3.B.3	Formal results for estimation and inference	. 140
		3.B.4	Proof of Proposition 3.1	. 142
		3.B.5	Proof of Proposition 3.2	. 144
		3.B.6	Proof of Proposition 3.3	. 148
		3.B.7	Special case of independence between K and ρ	. 150
		3.B.8	Naive method and small-unit bias	. 152
	App	endix 3.	C Supplements to the application	. 155
		3.0.1	Additional former	. 150
	App	3.0.2	D Formal link with the model of Contrikow et al. (2010)	. 109
	Арр	3 D 1	CST's framework	. 105
		3 D 2	Connection between GST's and our model	166
		0.12.2		. 100
4	Nor	asymp	totic bounds for Edgeworth expansions	169
	4.1	Introdu	action	. 170
	4.2	Contro	l of $\Delta_{n,E}$ under moment conditions only $\ldots \ldots \ldots \ldots \ldots \ldots$. 173
	4.3	Improv	ed bounds on $\Delta_{n,E}$ under assumptions on the tail behavior of f_{S_n}	. 175
	4.4 Conclusion and statistical applications			
	App	endix 4.	A Proof of the main results	. 188
		4.A.1	A smoothing inequality	. 188
		4.A.2	Outline of the proofs of Theorems 4.1 and 4.2	. 189
		4.A.3	Proof of Theorem 4.1 under Assumption 4.1	. 191
		4.A.4	Proof of Theorem 4.1 under Assumption 4.2	. 199

4.A.5 Proof of Theorem 4.2 under Assumption 4.1	202
4.A.6 Proof of Theorem 4.2 under Assumption 4.2	205
Appendix 4.B Technical lemmas	208
4.B.1 Control of the term Ω_1	208
4.B.2 Control of the residual term in an Edgeworth expansion under	
Assumption 4.1	209
4.B.3 Control of the residual term in an Edgeworth expansion under	
Assumption 4.2	211
4.B.4 Two bounds on incomplete Gamma-like integrals	213
4.B.5 Proof of Proposition 4.4	214
	015
5 Nonasymptotic confidence intervals in linear models	210
5.1 Introduction	210
5.2 Quality measures for confidence sets	220
5.5 Linear regression without endogeneity	220 000
5.5.1 Model and standard asymptotic inference	220
5.5.2 Our confidence interval	220
5.4. Dractical considerations	229
5.4 Flactical considerations	229
5.5 Simulations	200 222
5.6 1 Model and standard Anderson Public inference	∠ <u></u> <u> </u> <u> </u> 233
5.6.2 Our confidence set	200
5.6.3 Theoretical results	230
Appendix 5 A Proof of results in Section 5.3	230
Appendix J.A. 1 1001 01 results in Section 3.5 \therefore 5.4.1 Nonsymptotic conservativeness for the Edgeworth regime	
5 A 2 Nonasymptotic conservativeness for the Exponential regime	240
5 A 3 Proof of Theorem 5.1	241 245
5 A 4 Proof of Theorem 5.2	246
Appendix 5 B Proof of results in Section 5.6	247
5.B.1 Nonasymptotic conservativeness for the Edgeworth and Exponential	- 11
regimes	247
5.B.2 Proof of Theorem 5.3	248
5.B.3 Proof of Theorem 5.4	248
Appendix 5.C Additional lemmas	249
Appendix 5.D Definition of the remainders of the bound on Edgeworth expansion	253
5.D.1 Definition of r_{e}^{gen}	253
5.D.2 Definition of r_n^{icont}	254
Appendix 5.E Additional figures \ldots	255

Summary in English

This thesis is divided into two independent parts.

The first three chapters deal with measures of segregation or polarization. The notion of segregation/polarization applies to various situations, but the formal modeling remains the same. Suppose that a population of individuals, comprised of a minority and a majority group, is allocated into units or makes choices over a set of options. In practice, units can be neighborhoods, firms, school classes, and the minority group might be immigrants versus natives when studying residential, occupational, or school segregation. The modeling also encompasses speech polarization, for instance. The US congresspeople are divided into Democrats and Republicans; the units or options are the items of a dictionary of words or phrases, and the speakers choose which words they use.

Qualitatively, there is some segregation or polarization whenever the allocation or choice process leads to a situation where the two groups tend to select distinct units/options. For example, in residential segregation, the minority individuals are concentrated in some units more than in others instead of being uniformly allocated. Regarding speech polarization, the presence of polarization means that Democrats and Republicans tend to use different words or phrases when they speak in Congress. Despite that intuitive notion, quantitative measures of segregation/polarization struggle with the so-called "small-unit bias."

Imagine a simple case to convey the idea. We consider a population of employees, half women, half men, and we study occupational segregation between gender: do women and men tend to work in the same workplaces (which are here the units) or, on the contrary, do they concentrate into distinct workplaces? Assume that there is no segregation in the sense that, for each workplace, the probability that an arbitrary employee working there is a woman is equal to one-half. In other words, women and men are allocated randomly over workplaces. Yet, as long as the sizes (that is, the number of individuals) of workplaces are small, the observed proportions of women and men will not be precisely equal to the theoretical 50-50% repartition owing to statistical fluctuations. In an extreme case with two employees per workplace, although women and men are allocated randomly, on average, one-fourth of the workplaces will have only women and another one-fourth only men. Consequently, in such small-unit situations, the traditional segregation indices that rely on the dispersion of the empirical proportions are upward-biased and wrongly point to a positive amount of segregation. Furthermore, the bias impedes reliable comparisons of segregation measures across countries or over time.

Such situations of small units are pervasive in applied research. For workplace and school segregation, a large share of firms have less than ten employees, and classes usually have between twenty and forty students. The bias also arises when the units are not small *per se*, but only surveys of individuals are available. For instance, if one studies residential segregation using data from the French Labor Force Survey (INSEE's *Enquête Emploi*), a unit is made of around fifty individuals.

The first chapter presents a Stata command (Stata is a widespread software to conduct statistical and econometric analyses) that implements three methods to estimate segregation indices robust to "the small-unit bias."

The second chapter applies these methods to quantify residential segregation in France between 1968 and 2019 along various dimensions (nationality, social status, labor market position, proxy of races). In this study, units are defined as clusters of around thirty adjacent housing, which are basic bricks in the sampling scheme of the French Labor Force Survey.

Finally, the third chapter studies speech polarization in the US Congress between 1873 and 2016 using transcripts of congressional debates.

The second part of this thesis is concerned with constructing nonasymptotic confidence intervals (CIs) for the individual coefficients of a linear regression model. Linear regression models are a ubiquitous method of econometric analyses.

The CIs rely on explicit upper bounds on the uniform distance between the cumulative distribution function of a standardized sum of independent centered random variables with moments of order four and its first-order Edgeworth expansion. These bounds are derived in the fourth chapter, which is more technical and closer to statistics and probability than the other chapters.

The last and fifth chapter uses these results to construct CIs that, at the same time, are (i) valid for any sample size (ii) without assuming parametric assumption such as the normality of error terms or independence between covariates and error terms (hence, our CIs allow for heteroskedasticity), (iii) have a closed-form expression, and (iv) whose length is asymptotically the same as the usual CI based on the t-statistic; thus our CIs have a coverage equals to the desired nominal level in the limit when the sample size goes to infinity.

Note

Although connected, the five chapters of this thesis are independent research articles. That is why this thesis presents some redundancies.

The title of this thesis is "contributions to measures of segregation or polarization and to nonasymptotic inference in linear models" in English, and "contributions aux mesures de ségrégation ou de polarisation et à l'inférence non-asymptotique dans les modèles linéaires" in French. The thesis is mainly written in English.

The first and third chapters are co-authored with Xavier D'Haultfœuille (CREST, ENSAE Paris, Institut Polytechnique de Paris), my thesis supervisor, and Roland Rathelot (CREST, ENSAE Paris, Institut Polytechnique de Paris; previously Warwick University).

The fourth and fifth chapters are co-authored with Alexis Derumigny (Delft University of Technology) and Yannick Guyonvarch (INRAE; previously Telecom Paris, Institut Polytechnique de Paris), two former Ph.D. students at CREST.

I am the only author for the second chapter.

Any opinions expressed in this thesis are those of the author and not of any institution. All possible errors remain my own.

Résumé substantiel en français

Cette thèse se compose de deux parties indépendantes.

Les trois premiers chapitres s'intéressent aux mesures de ségrégation ou de polarisation. Ces notions s'appliquent à des contextes variés mais partagent une formalisation commune. Une population d'intérêt est divisée en deux groupes exclusifs, un groupe dit minoritaire et un groupe majoritaire, et les individus de ces deux groupes se répartissent entre des unités ou choisissent entre des options. Les unités sont par exemple des aires résidentielles, le groupe minoritaire les individus de nationalité étrangère et le groupe majoritaire les individus de nationalité française pour étudier un aspect de la ségrégation résidentielle en France. Cette modélisation permet également d'étudier la polarisation du langage politique aux États-Unis en considérant comme unités ou options les entrées d'un dictionnaire et comme groupes minoritaire et majoritaire les parlementaires du parti républicain et du parti démocrate ; ici, les individus choisissent les mots qu'ils utilisent.

Qualitativement, il y a de la ségrégation ou de la polarisation lorsque les deux groupes tendent à choisir systématiquement des options distinctes. En ségrégation résidentielle, cela signifie que les individus du groupe minoritaire sont concentrés dans certaines aires géographiques au lieu d'être répartis uniformément sur le territoire. Dans l'étude du langage politique, ce sera le cas lorsque certains mots ou expressions seront davantage prononcés par un parti que par un autre. Quantitativement, la mesure de la magnitude de la ségrégation ou polarisation est confrontée au *small-unit bias* (biais de petites unités).

Considérons un cas simple pour expliquer ce biais. Dans une population d'employés comprenant moitié d'hommes et moitié de femmes, on s'intéresse à la ségrégation professionnelle homme-femme : est-ce que les hommes et femmes travaillent dans les mêmes entreprises (qui jouent ici le rôle des unités) ou à se concentrer dans des entreprises distinctes ? Imaginons qu'il n'y ait pas de ségrégation au sens où, pour chaque entreprise, la probabilité qu'un employé quelconque soit une femme vaut un demi : les hommes et les femmes se répartissent aléatoirement entre les entreprises. Pour autant, si les entreprises sont de petites tailles, on n'observa pas des proportions hommes-femme égales exactement à 50-50. En raison des fluctuations statistiques, les proportions empiriques s'éloigneront du 50-50 théorique en probabilités. À l'extrême, pour des entreprises de deux employés, en moyenne, un quart d'entre elles n'auront que des femmes et un autre quart seulement des hommes. Ainsi, dans ces situations de petites unités, les mesures traditionnelles de ségrégation fondées sur la dispersion des proportions empiriques reporteront à tort une ségrégation élevée. De surcroît, en raison de ce biais, la comparaison des indices entre différents pays ou sur plusieurs périodes est douteuse.

Ces situations interviennent dans de nombreuses applications pratiques : une grande majorité d'entreprises a moins de dix salariés ; les classes comportent typiquement entre une vingtaine et trentaine d'élèves ; dans certaines données permettant d'étudier la ségrégation résidentielle, comme les grappes de l'Enquête Emploi de l'INSEE (Institut National de la Statistique et des Etudes Economiques), les unités regroupent une cinquantaine de personnes résidant dans des logements adjacents.

Le premier chapitre présente un programme Stata (logiciel courant de statistiques et d'économétrie) qui implémente trois méthodes permettant d'estimer des indices de ségrégation robustes au « small-unit bias ».

À partir des données de l'Enquête Emploi permettant de construire des unités formées d'une trentaine de logements adjacents, le second chapitre applique ces méthodes pour quantifier la ségrégation résidentielle en France entre 1968 et 2019 sur plusieurs dimensions (nationalité, statut social, position sur le marché du travail et proxy de l'ethnicité).

Le troisième chapitre utilise les retranscriptions des débats au Congrès américain entre 1873 et 2016 pour étudier la polarisation du langage politique.

La deuxième partie s'intéresse à la construction d'intervalles de confiance (IC) nonasymptotiques pour les coefficients des modèles de régression linéaire, un outil classique d'analyse économétrique.

Ces intervalles reposent sur des majorations explicites de la distance entre la distribution empirique d'une somme normalisée de variables aléatoires indépendantes centrées admettant des moments d'ordre quatre et son expansion d'Edgeworth de 1er ordre. Ces majorations sont obtenues dans le quatrième chapitre, plus statistique et technique.

Le cinquième et dernier chapitre utilise ces dernières pour construire des IC qui sont simultanément (i) valides pour toute taille d'échantillon (ii) sans imposer une distribution paramétrique de type Gaussienne ou l'indépendance entre régresseurs et résidus (et autorisent donc des résidus hétéroscédastiques), (iii) ayant une expression explicite, et (iv) atteignant asymptotiquement la même précision que les IC usuels fondés sur la normalité asymptotique de la statistique de Student.

Remarque

Malgré certains liens, les cinq chapitres constituant cette thèse sont des articles de recherche indépendants. De ce fait, ce manuscrit peut présenter quelques répétitions.

Le titre de cette thèse est "contributions aux mesures de ségrégation ou de polarisation et à l'inférence non-asymptotique dans les modèles linéaires" en français et "contributions to measures of segregation or polarization and to nonasymptotic inference in linear models" en anglais, langue de rédaction principale de ce manuscrit.

Le premier et le troisième chapitres sont coécrits avec mon directeur de thèse, Xavier D'Haultfœuille (CREST, ENSAE Paris, Institut Polytechnique de Paris), et Roland Rathelot (CREST, ENSAE Paris, Institut Polytechnique de Paris).

Le quatrième et le cinquième chapitres sont coécrits avec Alexis Derumigny (Delft University of Technology) et Yannick Guyonvarch (INRAE), deux anciens doctorants au CREST.

Je suis l'unique auteur du second chapitre.

Les opinions exprimées dans cette thèse n'engagent que leur auteur, qui assume l'entière responsabilité des erreurs qui pourraient subsister.

Introduction in English

This Ph.D. thesis in econometrics consists of two independent parts. It would be somewhat artificial and more of a rhetorical exercise to gather those two parts into a single title or topic. A first try would be "small-unit and small-sample issues in econometrics," because the two parts are concerned about something small. However, such a presentation might be misleading and, eventually, counterproductive because the object that is small differs in the two settings in the sense that the asymptotics are different.

This introduction briefly presents the two parts of this thesis and the related chapters, hence specifying that difference.

Each chapter begins with some context information and a summary of the chapter in the form of an abstract; I do not replicate those summaries in this general introduction.

A large sample of small units or a small sample of observations?

The first three chapters deal with the measure of segregation or polarization. In an abstract formulation that subsumes various applied situations, the problem is the following. In a two-group setting (extensions to multi-group settings are possible but not trivial), imagine you have a population of interest split into two exogenous groups: a minority group versus a majority one. Each individual of the population makes choices among a set of options, at least one choice, possibly several.¹

The groups are said exogenous in so far as group membership does not depend on the choices made. Group membership is an individual characteristic determined ex ante. It would be another problem not addressed in this thesis to construct coherent groups from the choices made (unsupervised clustering or supervised classification).

The terms "option" and "choice" take concrete meaning depending on the application. In Chapters 1 and 2, the context is that of residential segregation: an option (sometimes also referred to as a "unit" in such settings of residential, school or occupational segregation) is a neighborhood of about 30 adjacent housing; the choices are residential location choices, individuals choose where they live. The application of Chapter 3 concerns speech polarization. In this setting, an option is a word (or a phrase) in a given dictionary, an entry of the dictionary. The choices are occurrences of those words; the individuals choose which words they use when they talk. In a marketing or empirical industrial organization setting, you may consider two distinct consumer groups and study how different their purchase behaviors are; an option is a possible product, a choice is a purchase. Figure 1 below presents a small sketch of the modeling.

¹In this thesis, "polarization" and "segregation" are mainly used as synonyms. If I were to formally distinguish the two notions, "segregation" would be reserved for situations where each individual in the studied population makes only *one* choice among the set of options; in contrast, "polarization" would be more general, allowing for multiple choices made by an individual.



Figure 1: Setting for measures of segregation or polarization.

Qualitatively, segregation or polarization arises whenever the choice or allocation process is such that the two groups tend to systematically choose different options. For instance, in speech polarization in the U.S. Congress, polarization means that Republican and Democrats speakers use different words when they talk. In the context of segregation, it is the case when there is a concentration of minority individuals in some units more than in others.²

The problem of measurement is to quantify, by a bounded number, called an index of segregation/polarization, the extent to which the members of the two groups choose different options.

Another interesting way to assess polarization relates to information and prediction: if one observes the option chosen by an individual, does it help to predict the group membership of the individual? For instance, if I observe that a speaker said the phrase "undocumented workers" during a Congressional debate, can I use this information to decide whether the speaker is Democrat or Republican? Imagine the minority or reference group is made of individuals with at least a college education; if I know that someone lives in central Paris, does it change my beliefs about whether she or he holds a college degree?

Let introduce some notation to go further.³ Let denote \mathcal{J} the set of options; they are J distinct options indexed by j. We observe a sample of n choices, that is (minimally), for each option j, we observe the number K_j^R (K_j^D) of choices made by members of the reference group R (of the other group D) that select option j; let $K_j := K_j^R + K_j^D$ denote

 $^{^{2}}$ Remark that, in line with the existing literature, the term "segregation" is used in this thesis in a neutral sense to refer to such concentration phenomenon, independently of the underlying causes.

³The notation used here is that of Chapter 3. This chapter presents the problem of measuring polarization/segregation in a more general way compared to Chapters 1 and 2 that rather focus on the context of (residential) segregation. In the latter, the notation X (respectively p) replaces the notation K^R (respectively ρ).

the total number of choices selecting option j that we observe.

A natural way to quantify polarization or segregation is to consider the variation across the empirical proportions $(K_j^R/K_j)_{j=1,...,J}$. Intuitively, the more variation across those proportions, the higher polarization. Indeed, if a large part of those proportions is close to either 0 or 1, it means that most options are chosen by only one of the two groups; that is, the two groups tend to choose distinct options. On the contrary, if there is little variation in the proportions $(K_j^R/K_j)_j$, the choices of the two groups are more or less identically distributed over the possible options. Natural polarization indices are thus defined as functions mapping $(K_j^R/K_j)_{j=1,...,J}$ to a scalar index, the more variation in the proportion, the higher the index.

Although intuitive, considering the variations in the observed proportions is problematic when the K_j are small. To explain this, let first introduce another object: for each option j, ρ_j is the conditional probability that, knowing the choice selects option j, the choice is made by a member of the reference group R (and not one of the other group D).⁴

When the objective is to measure the level of polarization, it is arguable that the correct benchmark of null polarization (for which the index is minimal, typically normalized to 0) is the randomness benchmark defined as the equality of all the probabilities $(\rho_j)_{j=1,...,J}$. This benchmark differs from the evenness benchmark characterized by the equality of all the proportions $(K_j^R/K_j)_{j=1,...,J}$. Throughout this thesis, we adopt the randomness benchmark to define and measure polarization: we say that there is some polarization if and only if the probabilities $(\rho_j)_{j=1,...,J}$ are not all equal. In other words, we define polarization as a function of the underlying probabilities ρ_j instead of the observed proportions K_i^R/K_j .⁵

We model $(K_j^R, K_j, \rho_j)_{j=1,...,J}$ as independent and identically distributed (i.i.d.) random variables. Therefore, we sometimes omit the subscript j to lighten notation: K^R, K, ρ is a generic instance of the variables, with the same distribution as K_j^R, K_j, ρ_j) for any j. Under this assumption, the equality of the probabilities ρ_j means that the distribution of ρ is a Dirac (degenerate distribution): ρ is a constant random variable.

 K_j^R/K_j , as the observed proportion of choices of option j made by group R, is an estimator of ρ_j ; however, it requires K_j to tend to infinity to be consistent. When K is small, the naive indices of polarization based on the observed proportions $(K_j^R/K_j)_{j=1,...,J}$ over-estimate the level of segregation; in the segregation literature, this issue is known as the small-unit bias. The idea is simply that, even if the $(\rho_j)_{j=1,...,J}$ are all equal, with finite K_j , there will be variation in the proportions $(K_j^R/K_j)_{j=1,...,J}$. In other words, in addition to the variation in the probabilities $(\rho_j)_j$, which defines the real level of polarization, the variation in the proportions K_j^R/K_j incorporates small-sample variability. The smaller K, the more acute the bias.

Furthermore, the small-unit bias impedes reliable comparisons of polarization levels over time or across contexts since the magnitude of the bias might change. For instance, Chapter 2 studies residential segregation in France between 1968 and 2019 using data from

⁴Remark that, through Bayes's rule, it is the "reverse" conditional probability that the one considered in the above paragraph that presents polarization in terms of information and prediction.

⁵The expressions of "real" or "true" level of polarization that may be found in this thesis are used to refer to this definition of polarization and emphasize the distinction with a proportion-based approach. Note that it does not imply the evenness benchmark is irrelevant. For example, when studying the consequences or effects of segregation, it is probably more sensible to consider the *realization* of the choice/allocation process, thus the proportions K_j^R/K_j . However, when the objective is to measure the level or magnitude of polarization, I believe one should study the underlying data-generating process parameters that produce the realizations, namely here the probabilities ρ_j ; hence the randomness benchmark.

the Labor Force Survey. Over this period, due to the evolution of the survey scheme, the size of the neighborhoods used in the analysis has decreased: $K \approx 70$ on average during the seventies while $K \approx 30$ in 2003-2019. Consequently, the magnitude of the bias has been higher in recent years. All else equal, if the real polarization level remains unchanged, naive indices that neglect the bias would thus point to an increase of segregation over time; yet, such an evolution is only a statistical artifact when measured with proportion-based indices. Measures robust to the small-unit bias are essential to accurately study temporal variations or compare different settings.

This bias arises in many applications since it is often the case that each option contains few choices, namely that K is small. In segregation contexts, a large share of firms (a fortiori, workplaces) have less than ten employees (occupational segregation), classes typically have between twenty and forty pupils (school segregation). In more general polarization contexts, the bias occurs whenever the size J of the choice set is large relative to the number of observed choices $n := \sum_{i=1}^{J} K_{j}$.

The crucial point to understand is that having more data will not solve the small-unit bias; it is not a matter of too small samples in that sense. Indeed, in school segregation, for instance, a larger sample means observing additional classes, namely increase J. Yet, it does not modify the fact that there is a limited number K of pupils per class, which remains finite and small enough in practice for the small-unit bias to matter. More generally, as long as the concrete meaning of the options (depending on the specific application) is such that their capacity, understood as the number K of choices they can receive, is limited, the small-unit bias arises. Formally, the asymptotic is in the dimension J tending to infinity (and thus n too, as a side effect) instead of a set-up with a fixed number J of options and a number n of observed choices going to infinity.

In fact, such an asymptotic with J going to infinity is not limited to segregation contexts with units (options) whose capacity is constrained. In various settings, observing more data in the sense of observing more choices (n increases) also implies observing more options (J = J(n) increases too). For instance, it is the case of speech or text data analyzed in Chapter 3. Indeed, an empirical relationship exists between the length n of a corpus of texts (number of choices, namely occurrences of words here) and the richness, the size J of its vocabulary (that is, the number of distinct words that appear in the corpus).⁶ As you observe more text, you also observe new words. Likewise, as you observe more purchase decisions of, say, cars, you also observe new types of cars in the sense that, as a researcher studying car purchase behavior, having more data allows you to consider thinner, more precisely defined products (which play the role of the options here). In those settings, an asymptotic representing more data by n tending to infinity while J is kept fixed arguably misrepresents the situation by escaping the relationship between Jand n. A more credible asymptotic is having both J and n tending together to infinity to represent those types of large or high-dimensional choice sets.

In a nutshell, the first part of this thesis deals with measuring polarization in such large choice sets, with an asymptotic in the number J of options going to infinity (and n too) whereas the number K of choices observed per option remains limited.

Chapter 1 presents a Stata command that implements three methods proposed in the segregation literature to address the small-unit bias.

Chapter 2 is more applied; it uses this Stata command to measure residential segregation in France between 1968 and 2019 with different definitions of minority and majority groups

 $^{^6\}mathrm{We}$ refer to Remark 3.1 in Chapter 3 for further details.

based on demographic and socioeconomic variables.

Chapter 3 proposes a statistical model and a method to identify and estimate a polarization index in large choice sets; it then applies this methodology to text data by studying speech polarization in the U.S. Congress between 1873 and 2016.

The second part of this thesis relates more to statistics and theoretical econometrics. Compared to the previous setting where the data contains two dimensions (the number of choices and the number of options), the data is standard: univariate observations X_1, \ldots, X_n with a sample size n in Chapter 4, and a n-sample of an outcome (real random variable) and covariates (real random vector) $(Y_1, X_1), \ldots, (Y_n, X_n)$ in Chapter 5 that studies linear regression models. In those cases, the only asymptotic that we consider is in the sample size n going to infinity. Yet, in contrast, the second part of this thesis is primarily concerned with nonasymptotic inference, that is, constructing statistical procedures whose guarantees hold for any sample size n instead of limit results when ntends to infinity.

Chapter 5 constructs nonasymptotic confidence intervals (CIs) for the individual coefficients of a linear regression model: CIs that have a probability of containing the true parameter at least equal to their nominal level for any sample size. Using the usual terminology of linear regressions, these CIs are derived both in exogenous cases and in endogenous settings with one endogenous regressor, one instrument, and additional exogenous covariates. The construction uses results derived in Chapter 4, namely nonasymptotic (valid for any sample size) explicit bounds on the uniform distance between the cumulative distribution function of a standardized sum of independent centered random variables and its first-order Edgeworth expansion. In order to present those notions and the main challenges of nonasymptotic CIs, the rest of the introduction presents the simple example of conducting inference on an expectation.⁷

Introduction to nonasymptotic inference: the example of an expectation

We observe $n \in \mathbb{N}^*$ independent and identically distributed (i.i.d.) real random variables $(D_i)_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}^D$. We denote by $\mathcal{D} \subseteq \mathbb{R}$ the support of the distribution of D_1 , denoted \mathbb{P}^D . $\mathcal{P}(\mathcal{D})$ denotes the set of all probability distributions on \mathcal{D} .

We assume the data follows some statistical model whose parameter can be divided into a finite-dimensional part and the remainder:

$$\mathbf{P}^{D} \in \left\{ \mathbf{P}_{\theta} \in \mathcal{P}(\mathcal{D}), \theta = (\theta_{1}, \theta_{2}) \in \Theta := \mathbb{R} \times \Theta_{2} \right\},\$$

where Θ_2 is a topological space. The expression captures typical semi-parametric models where one is interested in a finite-dimensional parameter θ_1 while, for the richness of the model, the data-generating process also depends on a possibly infinite-dimensional "nuisance" parameter θ_2 .

In this example, $\mathcal{D} := \mathbb{R}$, $\theta_1 := \int u \, d\mathbb{P}^D \in \mathbb{R}$ is simply the expectation of D_1 , $\theta_2 := \mathbb{P}^D$ is the distribution of the observations, and Θ_2 the set of possible distributions on \mathbb{R} with finite variance: $\Theta_2 := \{P \in \mathcal{P}(\mathbb{R}) : \int u^2 dP < +\infty\}.$

We seek to construct a confidence interval (CI) for θ_1 . Essentially, a confidence interval $CI(1 - \alpha, n)$ is a random interval of \mathbb{R} that depends on the data and comes with a desired

⁷Note that this example is inspired by Example 4.1 of Chapter 4 and Example 5.1 of Chapter 5; it directly uses several parts of Chapter 5.

nominal level $1 - \alpha \in (0, 1)$, whose objective is to contain the target parameter θ_1 with a probability at least $1 - \alpha$, ideally equal to $1 - \alpha$. Ideally equal to $1 - \alpha$ exactly because, thinking back to the original Neyman-Pearson principle for parametric tests (the distinction of Type 1 and Type 2 errors) and the connection with CIs, the choice of a nominal level $1 - \alpha$ means we are willing to accept that the CI misses the parameter with a probability at most α in exchange for increased precision, through smaller length of the confidence interval. In other words, it is not desirable that $CI(1 - \alpha, n)$ contains θ_1 with a probability higher than $1 - \alpha$ because, in such a case, it means we are losing some precision: for the same stated nominal level, we could have a CI with a smaller length.

Quality of a confidence interval Several criteria exist to assess the quality of a given confidence interval, $CI(1 - \alpha, n)$. Let Θ be a fixed subset of Θ .

From an asymptotic point of view, $CI(1 - \alpha, n)$ is said to be asymptotically exact pointwise over $\widetilde{\Theta}$ if

$$\forall \theta \in \widetilde{\Theta}, \lim_{n \to +\infty} \mathbb{P}_{\mathbb{P}_{\theta}^{\otimes n}} \left(\operatorname{CI}(1 - \alpha, n) \ni \theta_1 \right) = 1 - \alpha.$$
(1)

A stronger asymptotic criterion exists: $CI(1 - \alpha, n)$ is said to be asymptotically exact uniformly over $\tilde{\Theta}$ if

$$\lim_{n \to +\infty} \sup_{\theta \in \widetilde{\Theta}} \left| \mathbb{P}_{\mathbf{P}_{\theta}^{\otimes n}} \Big(\mathrm{CI}(1-\alpha, n) \ni \theta_1 \Big) - (1-\alpha) \right| = 0.$$
⁽²⁾

If $\liminf_{n \to +\infty} \inf_{\theta \in \widetilde{\Theta}} \mathbb{P}_{P_{\theta}^{\otimes n}} \left(\operatorname{CI}(1 - \alpha, n) \ni \theta_1 \right) \ge 1 - \alpha$, the CI is said to be asymptotically conservative uniformly over $\widetilde{\Theta}$.

From a nonasymptotic perspective, we say that $CI(1 - \alpha, n)$ is nonasymptotically conservative over $\widetilde{\Theta}$ if

$$\forall n \ge 1, \inf_{\theta \in \widetilde{\Theta}} \mathbb{P}_{\mathsf{P}_{\theta}^{\otimes n}} \left(\operatorname{CI}(1 - \alpha, n) \ni \theta_1 \right) \ge 1 - \alpha.$$
(3)

This property evolves into nonasymptotic exactness over Θ of $\operatorname{CI}(1-\alpha, n)$ if (3) holds and, in addition, for any $n \geq 1$, $\sup_{\theta \in \widetilde{\Theta}} \mathbb{P}_{\mathrm{P}_{\theta}^{\otimes n}} \left(\operatorname{CI}(1-\alpha, n) \ni \theta_1 \right) = 1-\alpha$.

To present these notions and the challenges of nonasymptotic inference, we construct below several confidence intervals for θ_1 . To this end, let introduce the empirical mean and the empirical variance of the observations:⁸

$$\overline{D}_n := \frac{1}{n} \sum_{i=1}^n D_i, \quad \widehat{\sigma}_n^2 := \frac{1}{n-1} \sum_{i=1}^n \left(D_i - \overline{D}_n \right)^2.$$

Asymptotic approach First, we can construct a confidence interval based on asymptotic results. $\hat{\sigma}_n^2$ is a consistent estimator of the variance of D by the Law of Large Numbers and the Continuous Mapping Theorem (CMT). Hence, the Central Limit Theorem, combined with Slutsky's lemma and the CMT, gives

$$\frac{\sqrt{n}}{\sqrt{\widehat{\sigma}_n^2}} \left(\overline{D}_n - \theta_1 \right) \stackrel{d}{\longrightarrow} \mathcal{N}(0, 1).$$

⁸We use the unbiased empirical variance, with n-1 in the denominator, to work as such for Cochran/Fisher's theorem later; it is asymptotically equivalent to have n instead of n-1.

From this, denoting $q_P(\tau)$ the quantile at order $\tau \in (0, 1)$ of a distribution P, we obtain that

$$\operatorname{CI}_{as}(1-\alpha,n) := \left[\overline{D}_n \pm q_{\mathcal{N}(0,1)}(1-\alpha/2)\frac{\sqrt{\widehat{\sigma}_n^2}}{\sqrt{n}}\right]$$

is asymptotically exact pointwise over Θ . More generally, such construction is often implicitly what (applied) econometricians or statisticians have in mind when they rely on the asymptotic normality of an estimator to conduct inference.

However, some impossibility results implies $CI_{as}(1-\alpha, n)$ cannot be nonasymptotically conservative over the whole parameter space Θ .⁹ More assumptions are needed to build CIs with nonasymptotic guarantees. From there, my understanding is that there are two main ways to proceed:

- 1. parametric assumptions on the distribution of D can provide nonasymptotic exactness on a "small" (in the sense that it is parametric, finite-dimensional) subset of Θ ;
- 2. concentration inequalities that require known bounds on some moments or other features (for instance, the support) of P^D can yield nonasymptotic conservativeness on large(r) (nonparametric, infinite-dimensional) subsets of Θ .

The choice between those two possibilities can be interpreted as a trade-off between 1. the precision (exactness) of a CI and 2. the uniformity of its guarantees, the richness of the model the CI applies to. We present below one classic example of construction 1. with a parametric Gaussian assumption and two examples of construction 2.

Nonasymptotic exactness through parametric (Gaussian) assumption Let denote \mathcal{T}_k a Student distribution with k degrees of freedom. By Cochran/Fisher's theorem (i.i.d. Gaussian observations), we have, for any sample size (provided $n \geq 2$ to compute the empirical variance),

$$rac{\sqrt{n}}{\sqrt{\widehat{\sigma}_n^2}} \left(\overline{D}_n - \theta_1
ight) \sim \mathcal{T}_{n-1}$$

Consequently,

$$\operatorname{CI}_{N}(1-\alpha,n) := \left[\overline{D}_{n} \pm q_{\mathcal{T}_{n-1}}(1-\alpha/2)\frac{\sqrt{\widehat{\sigma}_{n}^{2}}}{\sqrt{n}}\right]$$

is nonasymptotically exact over $\widetilde{\Theta}_N \subsetneq \Theta$ with

$$\widetilde{\Theta}_N := \left\{ \theta \in \Theta : \mathcal{P}_{\theta} = \mathcal{N}(\theta_1, \sigma^2), \theta_1 \in \mathbb{R}, \sigma^2 \in \mathbb{R}_+^* \right\}.$$

 Θ_N can be deemed a "small" subset of Θ insofar as it is parametric. Besides, $\operatorname{CI}_N(1-\alpha, n)$ is asymptotically exact pointwise over the whole Θ as its asymptotic behavior is equivalent to that of $\operatorname{CI}_{as}(1-\alpha, n)$ when n goes to $+\infty$.

⁹Chapter 5 provides further details and references.

Nonasymptotic conservativeness through nonparametric assumptions (bounds on the variance or on the support) For some known $M < +\infty$, let define $\tilde{\Theta}_{BC} := \{\theta \in \Theta : \mathbb{V}_{P_{\theta}}[D] \leq M\}$. By Bienayme-Chebyshev (BC) inequality,

$$\operatorname{CI}_{BC}(1-\alpha,n) := \left[\overline{D}_n \pm \frac{\sqrt{M}}{\sqrt{\alpha n}}\right]$$

is nonasymptotically conservative over $\tilde{\Theta}_{BC} \subsetneq \Theta$. Compared to $\tilde{\Theta}_N$, $\tilde{\Theta}_{BC}$ is a large nonparametric subset of Θ .

For some known $-\infty < a < b < +\infty$, let define $\Theta_H := \{\theta \in \Theta : \text{support}(P_\theta) \subseteq [a, b]\}$. By Hoeffding (H) inequality,

$$\operatorname{CI}_{H}(1-\alpha,n) := \left[\overline{D}_{n} \pm \frac{(b-a)}{\sqrt{2}} \frac{\sqrt{\log(2/\alpha)}}{\sqrt{n}}\right]$$

is nonasymptotically conservative over $\tilde{\Theta}_H \subsetneq \Theta$. Likewise, as opposed to $\tilde{\Theta}_N$, $\tilde{\Theta}_H$ is large since nonparametric.

Nonasymptotic conservativeness and asymptotic exactness To compare the lengths of $\operatorname{CI}_{as}(1-\alpha, n)$, $\operatorname{CI}_{BC}(1-\alpha, n)$, and $\operatorname{CI}_H(1-\alpha, n)$, remark that, for any $\alpha \in (0, 1)$, $q_{\mathcal{N}(0,1)}(1-\alpha/2) < \sqrt{2\log(2/\alpha)}$ and $q_{\mathcal{N}(0,1)}(1-\alpha/2) < 1/\sqrt{\alpha}$. This implies that the CIs based on BC and H inequalities are asymptotically conservative. As explained above, being conservative is a negative feature of a CI, a fortiori asymptotically because it entails a loss of precision even when the sample size goes to infinity.

Nonasymptotic exactness is often quite demanding for it typically requires parametric assumptions as regards the distribution of the observations. Such assumptions might be restrictive in applications.

All in all, an interesting and feasible objective for nonasymptotic inference is the construction of CIs that are nonasymptotically conservative and, at the same time, asymptotically exact. To construct such CIs, the form of $\text{CI}_{as}(1-\alpha,n)$ suggests that something like $q_{\mathcal{N}(0,1)}(1-\alpha/2+\delta_n)$, for the α -related part of the CI's length, with adequate δ_n to ensure nonasymptotic conservativeness and $\lim_{n\to+\infty} \delta_n = 0$ to obtain asymptotic exactness might be interesting. This is exactly what enable to do Berry-Esseen inequality and Edgeworth expansion.

Berry-Esseen inequality and Edgeworth expansion This paragraph briefly presents Berry-Esseen (BE) inequality and Edgeworth expansion (EE).¹⁰ To begin with, it can be interesting to recall some asymptotic and nonasymptotic results to locate BE inequality among those.

The (weak) Law of Large Numbers (LLN) is an asymptotic result and states that $\overline{D}_n \xrightarrow{P} \mathbb{E}[D]$. Concentration inequalities (such as BC and H) are nonasymptotic results: for any sample size, they provide a number $t_{n,\alpha}$ such that the event $\{|\overline{D}_n - \mathbb{E}[D]| \leq t_{n,\alpha}\}$ holds with high probability (at least $1-\alpha$). They can be seen as a nonasymptotic equivalent or quantification of the LLN.

The Central Limit Theorem (CLT) yields an approximation of the distribution of \overline{D}_n when the sample size goes to infinity: informally, $\overline{D}_n \sim \mathcal{N}(\mathbb{E}[D], \mathbb{V}(D)/n)$ holds

¹⁰Chapter 4 provides more details.

asymptotically. In that sense, the CLT specifies the rate of convergence of the LLN, the usual $1/\sqrt{n}$.¹¹

BE inequality controls the uniform distance between the cumulative distribution function (c.d.f) of \overline{D}_n , properly centered and standardized, and Φ , the c.d.f of the standard Gaussian distribution $\mathcal{N}(0, 1)$:

$$\sup_{x \in \mathbb{R}} \left| \mathbb{P}\left(\sqrt{n} \, \frac{\overline{D}_n - \mathbb{E}[D]}{\sqrt{\sigma^2}} \le x \right) - \Phi(x) \right| \le \frac{1.88\mathbb{E}(|D - \mathbb{E}[D]]|^3)}{\sigma^3 \sqrt{n}},$$

with σ^2 the variance of D.¹² In that respect, BE inequality can be seen as a nonasymptotic equivalent or quantification of the CLT. Note that, compared to the CLT, BE inequality requires an additional moment: a finite third-order moment for D.

Let introduce additional notation to present EE. We use here the notation of Chapter 4. Let $(X_i)_{i=1,...,n}$ be a sequence of independent but not necessarily identically distributed (i.n.i.d.) real random variables that are assumed to be centered and with finite fourth-order moment: $\gamma_i := \mathbb{E}[X_i^4] < +\infty$. Let define

- the standard deviation of the sum of the X_i , $B_n := \sqrt{\sum_{i=1}^n \mathbb{E}[X_i^2]}$, so that the standardized sum of the X_i is $S_n := \sum_{i=1}^n X_i/B_n$,
- the average individual standard deviation $\overline{B}_n := B_n / \sqrt{n}$,
- the individual standard deviation $\sigma_i := \sqrt{\mathbb{E}[X_i^2]}$,
- the average standardized third raw moment $\lambda_{3,n} := \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}[X_i^3] / \overline{B}_n^3$
- the average standardized *p*-th absolute moment $K_{p,n} := \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}[|X_i|^p]/(\overline{B}_n)^p$, for $p \in \mathbb{N}^*$.

In the *i.i.d.* case, we can omit the subscript *i* for simplicity and we have $\sigma_i = \sigma$ the standard deviation of X, $B_n = \sigma \sqrt{n}$, $K_{4,n} = K_4 = \mathbb{E}[|X|^4]/\sigma^4$ the kurtosis of X (a measure of tail thickness) and $S_n = \sqrt{n}(\overline{X}_n - \mathbb{E}[X])/\sigma$.

Berry-Esseen inequalities aim at bounding the quantity

$$\Delta_{n,\mathrm{B}} := \sup_{x \in \mathbb{R}} \left| \mathbb{P} \left(S_n \le x \right) - \Phi(x) \right|.$$

Edgeworth expansions are a refinement of BE inequalities to adjust for the presence of nonasymptotic skewness in the distribution of S_n . They aim at controlling

$$\Delta_{n,\mathrm{E}} := \sup_{x \in \mathbb{R}} \left| \mathbb{P}\left(S_n \le x\right) - \Phi(x) - \underbrace{\frac{\lambda_{3,n}}{6\sqrt{n}}(1-x^2)\phi(x)}_{=:\mathrm{Edg}_n(x)} \right|,$$

where $\Phi(x) + \frac{\lambda_{3,n}}{6\sqrt{n}}(1-x^2)\phi(x)$ is called the one-term Edgeworth expansion of $\mathbb{P}(S_n \leq x)$ (ϕ is the density function of a $\mathcal{N}(0,1)$, the derivative of Φ).

¹¹Remark that concentration inequalities also specify the rate of convergence of the LLN.

¹²We express here BE inequality with the original constant 1.88 of Berry (1941). More recent results have derived lower constants. One of the contributions of Chapter 4 is to improve this inequality assuming finite fourth-order moments.

We can wonder why we are interested in bounds on the quantities $\Delta_{n,B}$ and $\Delta_{n,E}$. The answer is that they provide adequate concentration inequalities to obtain nonasymptotically conservative but asymptotically exact CIs. Indeed, for every $n \in \mathbb{N}^*$, for every t > 0, we have (disjoint union and complementary event)

$$\begin{split} \mathbb{P}\left(\frac{\sqrt{n}\left|\overline{X}_{n}\right|}{\sigma} > t\right) &\leq 1 - \mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq t\right) + \mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq -t\right) \\ \text{(add and subtract to see an EE)} &\leq 1 - \left(\mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq t\right) - \Phi(t) - \text{Edg}_{n}(t)\right) \\ &+ \left(\mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq -t\right) - \Phi(-t) - \text{Edg}_{n}(-t)\right) \\ &- \Phi(t) - \text{Edg}_{n}(t) + \Phi(-t) + \text{Edg}_{n}(-t) \\ \text{(triangular } \leq, \text{ parity of Edg}_{n}(\cdot), &\leq 2\Phi(-t) + \left|\mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq t\right) - \Phi(t) - \text{Edg}_{n}(t)\right| \\ &\text{and } \Phi(t) = 1 - \Phi(-t)) + \left|\mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq -t\right) - \Phi(-t) - \text{Edg}_{n}(-t)\right| \\ &\text{(definition of } \Delta_{n,\text{E}}) &\leq 2\left\{\Phi(-t) + \Delta_{n,\text{E}}\right\}. \end{split}$$

Finally, if we solve in t the equation

$$2\left\{\Phi(-t) + \Delta_{n,\mathrm{E}}\right\} = \alpha \iff \Phi(t) = 1 - \frac{\alpha}{2} + \Delta_{n,\mathrm{E}},$$

and let $X_i := D_i - \mathbb{E}[D]$, taking the complementary event of the previous computation yields

$$\Pr\left(\frac{\sqrt{n}\left|\overline{D}_{n}-\mathbb{E}[D]\right|}{\sigma} \le q_{\mathcal{N}(0,1)}\left(1-\frac{\alpha}{2}+\Delta_{n,\mathrm{E}}\right)\right) \ge 1-\alpha.$$
(4)

If we assume σ , the standard deviation of X (or equivalently of D) is known (the case of inference on an expectation with known variance), provided a bound δ_n on $\Delta_{n,E}$ such that $\lim_{n\to+\infty} \delta_n = 0$ is available, this gives the CI we have been looking for: (i) the result of Equation (4) holds for any sample size n, hence nonasymptotic conservativeness; (ii) its length converges to that of $\operatorname{CI}_{as}(1-\alpha, n)$, yielding asymptotic exactness.

When σ is unknown, an additional step is required to replace it by the estimator $\sqrt{\hat{\sigma}_n^2}$, but the idea is similar.

In essence, Chapter 5 proceeds the same analysis for the Ordinary Least Square (OLS) estimator of a linear regression model instead of the mean \overline{D}_n .

Asymptotic exactness is a desired theoretical property but, for practical use, for a fixed nominal level, the lower the bound δ_n on $\Delta_{n,E}$, the lower the length of the CI, hence better precision. This explains the interest of finding as small as possible constants for the bounds on $\Delta_{n,B}$ and $\Delta_{n,E}$. To our knowledge, the best existing bounds are due to Shevtsova (2013), with $\Delta_{n,B} \leq 0.5583 K_{3,n}/\sqrt{n}$ (*i.n.i.d.* case) and $\Delta_{n,B} \leq 0.4690 K_{3,n}/\sqrt{n}$ (*i.i.d.* case). They are derived assuming only finite third-order moments.

In linear regressions, the standard way to conduct inference, based on the asymptotic normality of the OLS estimator, requires finite fourth-order moments for the covariates. This is why Chapter 4 aims at deriving improved bounds under the assumption of finite fourth-order moments. We do so by obtaining bounds on $\Delta_{n,E}$, which implies bounds on $\Delta_{n,B}$. Indeed, the triangle inequality and $\sup_{x\in\mathbb{R}}(|1-x^2|\phi(x)/6) < 0.0665$ give

$$\Delta_{n,\mathrm{B}} \le \Delta_{n,\mathrm{E}} + \frac{0.0665\,\lambda_{3,n}}{\sqrt{n}}.$$

Finally, recall that the confidence intervals $\operatorname{CI}_{BC}(1-\alpha,n)$ (respectively $\operatorname{CI}_H(1-\alpha,n)$) requires a known bound M (known bounds a and b) on the variance (support) of the observations. Similarly, the above CI derived from Berry-Esseen inequality and Edgeworth expansions requires bounds on some moments to compute δ_n , namely on $\lambda_{3,n}$ (the skewness of X in the *i.i.d.* case), $K_{3,n}$ (the standardized absolute third moment), and $K_{4,n}$ (the kurtosis of X in the *i.i.d.* case). It happens that a bound on $K_{4,n}$ is in fact sufficient because $|\lambda_{3,n}| < 0.621K_{3,n}$ (Pinelis, 2011) and $K_{3,n} \leq K_{4,n}^{3/4}$ (Jensen inequality). Absent strong parametric assumption (like the Gaussian distribution above in this example), it appears difficult for nonasymptotic inference procedures to escape setting such bounds. In practice, this raises the question of how to choose them. Chapters 4 and 5 discuss how to it is possible to do so.

Introduction en français

Cette thèse en économétrie se compose de deux parties indépendantes. Réunir ces deux parties en une seule thématique ou un titre synthétique serait un exercice quelque peu artificiel, bien que possible dans la mesure où les deux parties ont en commun le fait de s'intéresser à de *petits* échantillons ou à de *petites* unités. Cela donnerait, par exemple, "Quelques problèmes économétriques du fait de petites unités ou de petits échantillons". Toutefois, cette présentation unifiée pourrait suggérer, à tort, une ressemblance entre les deux parties alors que la modélisation statistique (et notamment l'asymptotique) considérée diffère. Cette brève introduction présente les deux parties de cette thèse en cherchant à préciser cette différence.

Dans le corps du manuscrit, rédigé en anglais, chaque chapitre débute par quelques informations contextuelles et un résumé (abstract). Une traduction française de ces éléments est présentée à la fin de cette introduction pour chacun des cinq chapitres.

Un grand échantillon de petites unités ou un échantillon avec peu d'observations ?

Les trois premiers chapitres de cette thèse s'intéressent à la mesure de la ségrégation ou polarisation. Une formulation générale de ce problème, qui regroupe diverses applications pratiques, est la suivante. On considère une population d'intérêt divisée en deux groupes exogènes : un groupe dit minoritaire et un groupe dit majoritaire (l'extension à un nombre supérieur de groupes est possible mais non triviale). Chaque individu de la population fait des choix parmi un ensemble d'options, au moins un choix, possiblement plusieurs.¹³

Les groupes sont qualifiés d'exogènes au sens où le fait d'appartenir au groupe minoritaire ou au groupe majoritaire ne dépend pas des choix réalisés. C'est une caractérisation individuelle donnée a priori. Un problème différent serait de construire des groupes cohérents à partir de l'observation des choix réalisés (problème de clustering non-supervisé) ou de classifier les individus comme membres du groupe minoritaire ou du groupe majoritaire selon leurs choix (problème de classification supervisée). Cette problématique n'est pas abordée dans cette thèse.

Les termes abstraits d'option et de choix prennent leurs significations concrètes selon les applications. Les chapitres 1 et 2 s'intéressent à la ségrégation résidentielle. Dans ce cas, une option (aussi appelé une *unité* dans un cadre de ségrégation résidentielle, scolaire ou professionnelle) est un voisinage, une aire résidentielle d'une trentaine de logements adjacents tandis qu'un choix est un choix de localisation de la résidence principale, les

¹³Tout au long de cette thèse, les mots "polarisation" et "ségrégation" sont principalement employés comme synonymes. S'il fallait les distinguer formellement, on pourrait restreindre le terme "ségrégation" aux situations dans lesquelles les individus font un *seul* choix (par exemple, la localisation de sa résidence principale) tandis que le terme "polarisation" s'appliquerait lorsque les individus peuvent faire plus d'un choix (par exemple des choix de consommation où un individu peut acheter plusieurs produits).

individus choisissent où ils vivent. L'application du chapitre 3 porte sur la polarisation du langage politique. Dans ce contexte, une option est un mot (ou un groupe de mots, une expression) d'un certain dictionnaire. Une option est ainsi une entrée de ce dictionnaire tandis que les choix correspondent aux occurrences de ces mots, les individus choisissent les mots qu'ils utilisent lorsqu'ils parlent. Dans une problématique microéconomique plus classique, cette formalisation comprend également l'étude des choix de consommation de deux groupes de consommateurs : à quel point ces deux groupes ont-ils des comportements d'achat différents ? Dans ce cadre, une option est simplement un produit et un choix un achat d'un produit donné. La Figure 2 ci-dessous schématise cette modélisation commune de la mesure de la ségrégation ou polarisation.

Figure 2: Modélisation commune du problème de mesure de la ségrégation ou polarisation.



Qualitativement, on dira qu'il y a de la ségrégation ou de la polarisation lorsque la façon dont les individus font leurs choix (le processus générateur des données dans une approche statistique ou économétrique) est telle que les deux groupes choisissent des options différentes. Par exemple, dans le contexte de la polarisation du discours politique au Congrès américain, la présence de polarisation signifie que les Républicains et les Démocrates emploient des mots différents lorsqu'ils s'expriment au Congrès. Dans le cadre de la ségrégation résidentielle, il y a ségrégation lorsque les individus du groupe minoritaire sont concentrés dans certaines aires résidentielles, sur-représentés dans certaines et sous-représentés dans d'autres.¹⁴

La question de la mesure de la polarisation consiste à quantifier, au moyen d'un nombre normalisé (typiquement entre 0 et 1) appelé indice de ségrégation ou de polarisation, dans quelle mesure les membres des deux groupes tendent à choisir des options différentes.

 $^{^{14}}$ Dans la lignée de la littérature académique existante, le terme "ségrégation" est utilisé tout au long de cette thèse dans un sens neutre, descriptif, afin de faire référence à un tel phénomène de concentration, qu'elles que soient les causes sous-jacentes.

Une alternative intéressante pour appréhender la polarisation utilise les notions d'information et de prédiction : si j'observe l'option choisie par un individu, est-ce que cela m'aide à prédire le groupe, minoritaire ou majoritaire, auquel appartient cet individu ? Par exemple, si j'entends qu'un représentant utilise l'expression "undocumented workers" (travailleurs sans-papier) durant un débat parlementaire, est-ce que cela m'aide pour déterminer si ce représentant est Démocrate ou Républicain ? Dans un contexte de ségrégation résidentielle, supposons qu'on définisse le groupe minoritaire comme composé des individus ayant un diplôme universitaire de niveau Licence ou plus, le groupe majoritaire comme le complément. Dans ce cadre, si j'apprends qu'un individu habite à Paris intra-muros, est-ce que cela modifie ma croyance sur son appartenance au groupe minoritaire ou majoritaire ?

Introduisons quelques notations pour formaliser ce problème de mesure de la polarisation.¹⁵ Soient \mathcal{J} l'ensemble des options possibles et J le nombre d'options distinctes, indexées par j. On suppose qu'on observe un échantillon de n choix parmi cet ensemble d'options, c'est-à-dire qu'au minimum, on observe pour chaque option $j = 1, \ldots, J$, les variables suivantes : le nombre K_j^R (respectivement K_j^D) de choix faits par des membres du groupe minoritaire ou de référence R (respectivement de choix faits par l'autre groupe D) optant pour l'option j et le nombre total $K_j := K_j^R + K_j^D$ de choix optant pour l'option j.

Une façon naturelle de quantifier la polarisation ou la ségrégation est de s'intéresser aux variations entre les proportions empiriques $(K_j^R/K_j)_{j=1,\dots,J}$. Intuitivement, plus il y a de variation parmi ces proportions, plus la polarisation est importante. En effet, imaginons qu'une majorité de ces proportions sont proches de 0 ou proches de 1, cela signifie que la plupart des options sont sélectionnées par seulement un des deux groupes ; autrement dit, les deux groupes tendent à choisir systématiquement des options différentes. Inversement, s'il y a peu de variation parmi les proportions $(K_j^R/K_j)_j$, c'est-à-dire si les quantités K_j^R/K_j sont plus ou moins constantes inter j, alors nous sommes dans une situation de faible polarisation dans laquelle les choix des deux groupes se répartissent de façon à peu près identique entre les options possibles. Une première définition naturelle d'un indice de polarisation sera donc une fonction associant à $(K_j^R/K_j)_{j=1,\dots,J}$ un indice de façon telle que l'indice est d'autant plus élevé que les variations sont importantes parmi les proportions $(K_j^R/K_j)_j$.

Bien qu'intuitive, l'utilisation de la variation au sein des proportions $(K_j^R/K_j)_j$ pose problème lorsque les K_j sont petits. Afin d'expliquer cela, définissons un autre objet : pour chaque option j, ρ_j est la probabilité conditionnelle, sachant que le choix a sélectionné l'option j, que le choix a été fait par un individu du groupe de référence R (par opposition à avoir été fait par un individu de l'autre groupe D).¹⁶

Lorsqu'on cherche à mesurer le niveau de polarisation, on peut soutenir que le bon point de comparaison, correspondant à l'absence de polarisation (où l'indice de polarisation prend sa valeur minimale, typiquement normalisée à 0), est la situation d'une répartition aléatoire ("randomness benchmark") définie par l'égalité de toutes les probabilités $(\rho_j)_{j=1,...,J}$. Ce point de comparaison est distinct d'une situation d'égalité des proportions empiriques ("evenness benchmark") caractérisée par l'égalité de toutes les proportions $(K_i^R/K_j)_{j=1,...,J}$.

¹⁵Les notations utilisées ici reprennent celles du chapitre 3. Celui-ci s'intéresse au problème de mesure de la ségrégation ou polarisation d'une façon plus générale par rapport aux chapitres 1 et 2, qui se focalisent plutôt sur le cadre de ségrégation résidentielle. Dans ces deux chapitres, la notation X (respectivement p) correspond à la notation K^R (respectivement ρ) employée dans cette introduction.

¹⁶Par la formule de Bayes, il s'agit de la probabilité conditionnelle "inverse" de celle considérée dans le paragraphe précédent qui présente la polarisation en termes d'information et de prédiction.

Dans cette thèse, on adopte la comparaison à une répartition aléatoire pour définir et mesurer la polarisation : on dira qu'il y a un niveau non nul de polarisation si et seulement si les probabilités $(\rho_j)_j$ ne sont pas toutes égales. Autrement dit, on définit le niveau de polarisation comme une fonction des probabilités sous-jacentes ρ_j et non en fonction des proportions observées K_j^R/K_j .¹⁷

On modélise $(K_j^R, K_j, \rho_j)_{j=1,...,J}$ comme des variables aléatoires indépendantes et identiquement distribuées (i.i.d.). C'est pourquoi on omettra parfois l'indice j afin d'alléger les notations : K^R , K, ρ est une instance générique de ces variables, ayant la même distribution que K_j^R, K_j, ρ_j) quel que soit j. Sous cette hypothèse, l'égalité des probabilités ρ_j signifie que la distribution de ρ est dégénérée : ρ est une variable aléatoire constante (masse de Dirac).

 K_j^R/K_j , étant la proportion des choix portant sur l'option j réalisés par le groupe R, est un estimateur de ρ_j . Cependant, il n'est consistent que si K_j tend vers l'infini. Lorsque Kest petit, les indices "naïfs" de polarisation, ceux fondés sur les proportions observées $(K_j^R/K_j)_{j=1,\dots,J}$, sur-estiment le niveau de ségrégation ; dans la littérature, ce problème est connu sous le nom de "biais de petites unités" ("small-unit bias"). L'idée est la suivante : même si les probabilités $(\rho_j)_{j=1,\dots,J}$ sont toutes égales, avec des K_j finis, il y aura des variations dans les proportions $(K_j^R/K_j)_{j=1,\dots,J}$. En d'autres termes, en plus de la variation au sein des probabilités $(\rho_j)_j$, laquelle définit le niveau réel de polarisation, la variation au sein des proportions $(K_j^R/K_j)_j$ incorpore des fluctuations aléatoires dues à la finitude des K_j . Ce biais est d'autant plus marqué que K est petit.

De surcroît, ce biais de petites unités empêche des comparaisons fiables du niveau de polarisation au cours du temps ou entre pays ou situations puisque l'ampleur du biais peut changer.

Par exemple, le chapitre 2 étudie la ségrégation résidentielle en France entre 1968 et 2019 à l'aide des données de l'Enquête Emploi. Au cours de cette période, suite à des changements dans la confection de l'enquête, la taille des unités (les voisinages ou aires résidentielles ici) a diminué au cours du temps : $K \approx 70$ en moyenne durant les années 1970 mais $K \approx 30$ entre 2003 et 2019. En conséquence, l'ampleur du biais est plus important dans les années plus récentes. Toutes choses égales par ailleurs, si le niveau réel de ségrégation restait constant, les indices naïfs de ségrégation utilisant les proportions K_j^R/K_j , qui ne prennent pas en compte le biais, suggéreraient une hausse de la ségrégation au cours du temps. Ce résultat serait toutefois fallacieux, étant dû uniquement au changement de la magnitude du biais de petites unités qui impacte les indices naïfs. Des mesures de la ségrégation robustes à ce biais de petites unités sont ainsi cruciales pour étudier de façon fiable les évolutions temporelles du niveau de polarisation ou réaliser des comparaisons entre pays.

Ce biais de petites unités survient dans de nombreuses applications empiriques puisqu'il est courant que chaque option ne puisse contenir qu'un nombre limité de choix, c'est-à-dire

¹⁷Les termes de niveau réel ("real") ou de vrai ("true") niveau de polarisation qu'on peut trouver dans cette thèse sont employés pour faire référence à cette définition de la polarisation en insistant sur la distinction avec une approche fondée sur les proportions K_j^R/K_j . Pour autant, cela ne signifie pas que la comparaison à une situation d'égalité des proportions n'est pas intéressante. En particulier, si l'on étudie les conséquences de la ségrégation, il est probablement préférable d'utiliser les réalisations du processus générant les choix des individus, c'est-à-dire les proportions empiriques observées K_j^R/K_j . Toutefois, lorsque l'objectif est la mesure du niveau de polarisation, je pense qu'il est plus pertinent d'étudier les paramètres du processus générateur des données (les données étant ici les choix réalisés par les individus, et partant les proportions K_j^R/K_j), paramètres que sont justement les probabilités ρ_j ; de là, le choix de cette définition de la polarisation fondée sur les probabilités ρ_j .

que K soit petit. Dans des contextes de ségrégation, une part majoritaire des entreprises (a fortiori des lieux de travail) compte moins de dix employés (ségrégation professionnelle), les classes comportent en général entre vingt et quarante élèves (ségrégation scolaire). Dans des contextes plus généraux de polarisation, ce biais est présent lorsque le nombre J d'options pouvant être choisies est grand relativement au nombre de choix observés $n := \sum_{j=1}^{J} K_j$.

Il est essentiel de comprendre qu'avoir plus de données ne résout en rien ce problème de biais de petites unités ; ce n'est pas un problème dû à de petites tailles d'échantillon en ce sens. En effet, dans le contexte de ségrégation scolaire par exemple, avoir un plus grand échantillon signifie observer plus de classes, c'est-à-dire augmenter J. Or, cela ne change pas le nombre K limité d'élèves par classe, qui reste fini et suffisamment petit en pratique pour que le biais de petites unités ne soit pas négligeable. Plus généralement, ce biais intervient dès lors que la définition concrète d'une option (laquelle dépend de l'application spécifique considérée) est telle que sa capacité, au sens du nombre K de choix sélectionnant cette option, est limitée. Formellement, l'asymptotique du modèle statistique porte sur la dimension J tendant vers l'infini (et partant, n tend également vers l'infini) par opposition à une modélisation avec un nombre J fixé d'options et un nombre n de choix observés tendant vers l'infini.

Plus généralement, cette asymptotique où J tend vers l'infini n'est pas limitée à des contextes de ségrégation où les unités (les options) ont une capacité contrainte. Dans divers cas, observer davantage de données au sens d'observer davantage de choix (n augmente) implique également d'observer davantage d'options (J = J(n) augmente aussi). Par exemple, c'est le cas pour des données textuelles ou des retranscriptions de discours qui sont analysées au chapitre 3. En effet, il existe une relation empirique, un fait stylisé reliant la longueur n d'un corpus de textes (le nombre de choix, soit ici le nombre total d'occurrences de mots) et la richesse, la taille J du vocabulaire de ce corpus (le nombre de mots distincts employés dans le corpus).¹⁸ Lorsqu'on observe plus de textes, on observe aussi de nouveaux mots. De façon analogue, si, par exemple, on observe davantage de choix d'achats de voitures, un chercheur, statisticien ou économètre étudiant le comportement d'achat des voitures peut se permettre, ayant davantage de données, de raffiner sa définition des produits, autrement dit, de distinguer comme des produits distincts des voitures proches qui étaient auparavant considérées comme le même produit (la même option) avec peu de données. Dans ce genre de situations, on peut défendre qu'une asymptotique représentant le fait d'observer plus de données par n tendant vers l'infini alors que J est maintenu fixe dénature la situation en ne prenant pas en compte la relation existant entre J et n. Une approximation asymptotique plus crédible est d'avoir J et n tendant simultanément vers l'infini pour représenter ce type de données de choix dans un grand nombre d'options, dans de grands ensembles de choix ("large or high-dimensional choice sets").

La première partie de cette thèse cherche à mesure de façon fiable le niveau de polarisation dans de telles situations avec une asymptotique du nombre J d'options tendant vers l'infini (et n tendant aussi vers l'infini par conséquent) alors que le nombre K de choix observés par option reste limité.

Le chapitre 1 présente une commande Stata qui met en oeuvre trois méthodes proposées dans la littérature pour prendre en compte le biais de petites unités.

Le chapitre 2 est plus appliqué ; il utilise cette commande Stata pour mesurer le niveau de ségrégation résidentielle en France entre 1968 et 2019 en comparant différentes définitions des groupes minoritaires et majoritaires en fonction de diverses variables démographiques

¹⁸Voir la Remarque 3.1 du chapitre 3 pour davantage de détails.

et socio-économiques.

Le chapitre 3 propose un modèle statistique et une méthode pour identifier et estimer un indice de polarisation dans de grands ensembles de choix ; il applique ensuite cette méthodologie aux données textuelles en étudiant la polarisation du langage politique à partir des discours prononcés au Congrès américain entre 1873 et 2016.

La deuxième partie de cette thèse traite de statistiques et d'économétrie théorique. Par rapport au cadre de la première partie où les données comportent deux dimensions (le nombre de choix et le nombre d'options), les données sont ici plus standard : n observations univariées $(X_1, \ldots, X_n$ au chapitre 4 et un n-échantillon d'une variable de résultat (variable aléatoire réelle) et de régresseurs (vecteur aléatoire) $(Y_1, X_1), \ldots, (Y_n, X_n)$ au chapitre 5 qui étudie des modèles de régression linéaire. Dans ces deux cas, l'asymptotique considérée dans cette thèse est en la taille de l'échantillon n tendant vers l'infini. Pour autant, la deuxième partie de cette thèse s'intéresse principalement à l'inférence non-asymptotique, c'est-à-dire à la construction de procédures statistiques dont les garanties théoriques sont vérifiées pour toute taille d'échantillon n par opposition à des résultats limites valides lorsque n tend vers l'infini.

Le chapitre 5 construit des intervalles de confiances (ICs) non-asymptotiques pour les coefficients individuels d'un modèle de régression linéaire : ces ICs ont une probabilité de contenir le paramètre d'intérêt au moins égale à leur niveau nominal pour toute taille d'échantillon. Ces ICs sont obtenus pour des régresseurs exogènes et également dans des situations d'endogénéité avec un régresseur endogène, un instrument, et d'éventuelles variables de contrôle exogènes. Ils reposent sur des résultats obtenus dans le chapitre 4 : des bornes non-asymptotiques (valides pour toute taille d'échantillon) explicites sur la distance (en norme infinie ou norme uniforme) entre la fonction de répartition d'une somme standardisée de variables aléatoires indépendantes et centrées et son expansion d'Edgeworth au premier ordre. Pour présenter ces notions et les principaux défis des CIs non-asymptotiques, le reste de cette introduction présente un exemple : l'inférence sur une simple espérance.¹⁹

Introduction à l'inférence non-asymptotique : le cas d'une simple espérance

On observe $n \in \mathbb{N}^*$ variables aléatoires réelles indépendantes et identiquement distribuées $(i.i.d.) \ (D_i)_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}^D$. On note $\mathcal{D} \subseteq \mathbb{R}$ le support de la loi de D_1 , laquelle est notée \mathbb{P}^D . $\mathcal{P}(\mathcal{D})$ désigne l'ensemble des lois de probabilités sur \mathcal{D} .

On suppose que les données suivent un certain modèle statistique dont le paramètre peut être décomposé en une partie de dimension finie (ici de dimension 1 pour l'espérance) et la partie restante :

$$\mathbf{P}^{D} \in \left\{ \mathbf{P}_{\theta} \in \mathcal{P}(\mathcal{D}), \theta = (\theta_{1}, \theta_{2}) \in \Theta := \mathbb{R} \times \Theta_{2} \right\},\$$

où Θ_2 est un espace topologique. Cette modélisation correspond au cas classique de modèles semi-paramétriques où l'on est intéressé par un paramètre fini-dimensionnel θ_1 mais, pour la richesse du modèle, le processus générateur des données ("data-generating process") dépend également d'un paramètre de nuisance θ_2 , lequel est potentiellement de dimension infinie.

 $^{^{19}}$ Cet exemple est inspiré des exemples 4.1 du chapitre 4 et 5.1 du chapitre 5. Il reprend directement certains développements du chapitre 5.

Dans cet exemple, $\mathcal{D} := \mathbb{R}$, $\theta_1 := \int u \, d\mathbb{P}^D \in \mathbb{R}$ est simplement l'espérance de D_1 , $\theta_2 := \mathbb{P}^D$ est la loi des données, et Θ_2 est l'ensemble des lois de probabilités sur \mathbb{R} ayant une variance finie : $\Theta_2 := \{P \in \mathcal{P}(\mathbb{R}) : \int u^2 dP < +\infty\}.$

On cherche à construire un intervalle de confiance (IC) pour θ_1 . Un intervalle de confiance (confidence interval) $\operatorname{CI}(1-\alpha, n)$ est essentiellement un intervalle aléatoire de \mathbb{R} , qui dépend des données, et qui vient avec un niveau nominal souhaité $1-\alpha \in (0,1)$. Son objectif est de contenir le paramètre cible θ_1 avec une probabilité au moins égale à $1-\alpha$ et, idéalement, exactement égale à $1-\alpha$.

Idéalement égale à $1 - \alpha$ puisque, si l'on repense au principe de Neyman-Pearson pour les tests paramétriques (la distinction entre les erreurs de première et de seconde espèce) et leurs liens avec les ICs, choisir un niveau nominal $1 - \alpha$ signifie qu'on est prêt à ce que (et même on souhaite) que l'IC ne contienne pas le paramètre avec une probabilité au plus α en échange d'une meilleure précision, au moyen d'un intervalle de confiance ayant une longueur plus petite. Autrement dit, il n'est pas souhaitable que $CI(1 - \alpha, n)$ contienne θ_1 avec une probabilité strictement supérieure à $1 - \alpha$ puisque, dans un tel cas, cela signifie qu'on perd en précision : pour le même niveau nominal souhaité, on pourrait avoir un CI avec une plus petite longueur.

Qualité d'un intervalle de confiance Plusieurs critères existent pour juger de la qualité d'un intervalle de confiance $CI(1 - \alpha, n)$ donné. Soit $\tilde{\Theta}$ un sous-ensemble fixé de Θ .

Dans une perspective asymptotique, on dit que $CI(1 - \alpha, n)$ est asymptotiquement exact ponctuellement sur $\tilde{\Theta}$ lorsque

$$\forall \theta \in \widetilde{\Theta}, \lim_{n \to +\infty} \mathbb{P}_{\mathbb{P}_{\theta}^{\otimes n}} \left(\operatorname{CI}(1 - \alpha, n) \ni \theta_1 \right) = 1 - \alpha.$$
(5)

Il existe un critère asymptotique plus fort : $CI(1 - \alpha, n)$ est dit être asymptotiquement exact uniformément sur $\tilde{\Theta}$ lorsque

$$\lim_{n \to +\infty} \sup_{\theta \in \widetilde{\Theta}} \left| \mathbb{P}_{\mathbf{P}_{\theta}^{\otimes n}} \left(\operatorname{CI}(1 - \alpha, n) \ni \theta_1 \right) - (1 - \alpha) \right| = 0.$$
(6)

Si $\liminf_{n \to +\infty} \inf_{\theta \in \widetilde{\Theta}} \mathbb{P}_{\mathsf{P}_{\theta}^{\otimes n}} \left(\operatorname{CI}(1-\alpha, n) \ni \theta_1 \right) \geq 1-\alpha, \operatorname{CI}(1-\alpha, n) \text{ est dit être asymptotiquement conservateur uniformément sur } \widetilde{\Theta}.$

Dans une approche non-asymptotique, on dit que $CI(1-\alpha, n)$ est non-asymptotiquement conservateur sur $\tilde{\Theta}$ lorsque

$$\forall n \ge 1, \inf_{\theta \in \widetilde{\Theta}} \mathbb{P}_{\mathsf{P}_{\theta}^{\otimes n}} \left(\operatorname{CI}(1 - \alpha, n) \ni \theta_1 \right) \ge 1 - \alpha.$$
(7)

Cette propriété devient l'exactitude non-asymptotique sur $\widetilde{\Theta}$ de $\operatorname{CI}(1-\alpha, n)$ si on a (7) et, de plus, pour tout $n \ge 1$, $\sup_{\theta \in \widetilde{\Theta}} \mathbb{P}_{\mathrm{P}_{\theta}^{\otimes n}} \left(\operatorname{CI}(1-\alpha, n) \ni \theta_1 \right) = 1-\alpha$.

Pour présenter ces notions et les défis de l'inférence non-asymptotique, on construit ci-dessous plusieurs intervalles de confiances pour θ_1 . A cette fin, on introduit la moyenne empirique et la variance empirique des observations :²⁰

$$\overline{D}_n := \frac{1}{n} \sum_{i=1}^n D_i, \quad \widehat{\sigma}_n^2 := \frac{1}{n-1} \sum_{i=1}^n \left(D_i - \overline{D}_n \right)^2$$

²⁰On utilise la variance empirique non-biaisée, avec un n-1 au dénominateur, pour pouvoir appliquer directement le théorème de Cochran-Fisher par la suite ; asymptotiquement, c'est équivalent à avoir un estimateur de la variance avec n et non n-1 au dénominateur.

Approche asymptotique Tout d'abord, on peut construire un intervalle de confiance reposant sur des résultats asymptotiques. Par la Loi des Grands Nombres et le Continuous Mapping Theorem (CMT, préservation des différents modes de convergence de variables aléatoires par des fonctions continues), $\hat{\sigma}_n^2$ est un estimateur consistant de la variance de D. Ainsi, le Théorème Central Limite, combiné avec le lemme de Slutsky et le CMT, permet d'écrire

$$\frac{\sqrt{n}}{\sqrt{\widehat{\sigma}_n^2}} \left(\overline{D}_n - \theta_1 \right) \stackrel{d}{\longrightarrow} \mathcal{N}(0, 1).$$

De là, en notant $q_P(\tau)$ le quantile d'ordre $\tau \in (0, 1)$ d'une distribution P, on en déduit que

$$\operatorname{CI}_{as}(1-\alpha,n) := \left[\overline{D}_n \pm q_{\mathcal{N}(0,1)}(1-\alpha/2)\frac{\sqrt{\widehat{\sigma}_n^2}}{\sqrt{n}}\right]$$

est asymptotiquement exact ponctuellement sur Θ . Plus généralement, ce genre de construction est souvent ce qu'ont implicitement en tête les économètres et statisticiens (appliqués) lorsque l'inférence se fonde sur la normalité asymptotique de l'estimateur utilisé.

Cependant, certains résultats d'impossibilité impliquent que $\operatorname{CI}_{as}(1-\alpha,n)$ ne peut être non-asymptotiquement conservateur sur l'intégralité de l'ensemble des paramètres Θ .²¹ Des hypothèses supplémentaires sont requises pour construire des ICs ayant des garanties non-asymptotiques. D'après ma compréhension du problème, il existe pour cela deux façons principales de procéder :

- 1. des hypothèses paramétriques sur la distribution de D peuvent permettre d'obtenir l'exactitude non-asymptotique sur un "petit" sous-ensemble de Θ (au sens où ce sous-ensemble est paramétrique, de dimension finie) ;
- 2. des inégalités de concentration qui requièrent la connaissance de bornes sur certains moments ou d'autres caractéristiques (comme, par exemple, le support) de P^D peuvent donner des ICs non-asymptotiquement conservateurs sur de (plus) grands sous-ensembles de Θ (au sens où ils sont non-paramétriques, de dimension infinie).

Le choix entre ces deux possibilités peut s'interpréter comme un arbitrage entre 1. la précision (exactitude) d'un IC et 2. l'uniformité de ses garanties, la richesse du modèle auquel il s'applique. Un exemple classique d'une construction de type 1. utilisant une hypothèse paramétrique gaussienne et deux exemples de la méthode 2. sont présentés ci-dessous.

IC non-asymptotiquement exact au moyen d'une hypothèse paramétrique (gaussienne) Notions \mathcal{T}_k une loi de Student avec k degrés de liberté. D'après le théorème de Cochran-Fisher (observations i.i.d. gaussiennes), on a, pour toute taille d'échantillon (sous réserve d'avoir $n \geq 2$ pour calculer la variance empirique),

$$\frac{\sqrt{n}}{\sqrt{\hat{\sigma}_n^2}} \left(\overline{D}_n - \theta_1 \right) \sim \mathcal{T}_{n-1}.$$

²¹Le chapitre 5 fournit davantage de détails et des références.

Par conséquent,

$$\operatorname{CI}_{N}(1-\alpha,n) := \left[\overline{D}_{n} \pm q_{\mathcal{T}_{n-1}}(1-\alpha/2)\frac{\sqrt{\widehat{\sigma}_{n}^{2}}}{\sqrt{n}}\right]$$

est non-asymptotiquement exact sur $\Theta_N \subsetneq \Theta$ où

$$\widetilde{\Theta}_N := \left\{ \theta \in \Theta : \mathcal{P}_{\theta} = \mathcal{N}(\theta_1, \sigma^2), \theta_1 \in \mathbb{R}, \sigma^2 \in \mathbb{R}_+^* \right\}.$$

 Θ_N peut être décrit comme un "petit" sous-ensemble de Θ dans la mesure où il est paramétrique. Par ailleurs, $\operatorname{CI}_N(1-\alpha, n)$ est asymptotiquement exact ponctuellement sur tout Θ puisque son comportement asymptotique est équivalent à celui de $\operatorname{CI}_{as}(1-\alpha, n)$ lorsque n tend vers $+\infty$.

IC non-asymptotiquement conservateur au moyen d'hypothèses nonparamétriques (bornes sur la variance ou le support) Pour un certain $M < +\infty$ connu, posons $\tilde{\Theta}_{BC} := \{\theta \in \Theta : \mathbb{V}_{P_{\theta}}[D] \leq M\}$. Par l'inégalité de Bienaymé-Chebychev (BC),

$$\operatorname{CI}_{BC}(1-\alpha,n) := \left[\overline{D}_n \pm \frac{\sqrt{M}}{\sqrt{\alpha n}}\right]$$

est non-asymptotiquement conservateur sur $\tilde{\Theta}_{BC} \subsetneq \Theta$. Par rapport à $\tilde{\Theta}_N$, $\tilde{\Theta}_{BC}$ est un grand sous-ensemble non-paramétrique de Θ .

Pour certains $-\infty < a < b < +\infty$ connus, posons $\Theta_H := \{\theta \in \Theta : \operatorname{support}(P_\theta) \subseteq [a, b]\}$. Par l'inégalité de Hoeffding (H),

$$\operatorname{CI}_{H}(1-\alpha,n) := \left[\overline{D}_{n} \pm \frac{(b-a)}{\sqrt{2}} \frac{\sqrt{\log(2/\alpha)}}{\sqrt{n}}\right]$$

est non-asymptotiquement conservateur sur $\widetilde{\Theta}_H \subsetneq \Theta$. De même, par rapport à $\widetilde{\Theta}_N$, $\widetilde{\Theta}_H$ est grand au sens de non-paramétrique.

IC non-asymptotiquement conservateurs et asymptotiquement exacts Pour comparer les longueurs de $\operatorname{CI}_{as}(1-\alpha,n)$, $\operatorname{CI}_{BC}(1-\alpha,n)$ et $\operatorname{CI}_H(1-\alpha,n)$, on peut remarquer que, pour tout $\alpha \in (0,1)$, $q_{\mathcal{N}(0,1)}(1-\alpha/2) < \sqrt{2\log(2/\alpha)}$ et $q_{\mathcal{N}(0,1)}(1-\alpha/2) < 1/\sqrt{\alpha}$. Ceci implique que les ICs reposant sur les inégalités de BC et de H sont asymptotiquement conservateurs. Comme expliqué plus haut, être conservateur est un défaut d'un IC, a fortiori asymptotiquement puisqu'il entraîne une perte de précision et ceci même lorsque la taille de l'échantillon tend vers l'infini.

Dans le même temps, l'exactitude non-asymptotique est souvent très exigeante car, typiquement, elle requiert des hypothèses paramétriques sur la distribution des données. De telles hypothèses peuvent être restrictives dans des applications pratiques.

Au final, un objectif intéressant et a priori atteignable pour une inférence nonasymptotique est de construire des ICs qui sont à la fois non-asymptotiquement conservateurs et asymptotiquement exacts.

Pour obtenir de tels ICs, la forme de $\operatorname{CI}_{as}(1-\alpha,n)$ suggère qu'une expression telle que $q_{\mathcal{N}(0,1)}(1-\alpha/2+\delta_n)$, pour la partie de la longueur de l'IC dépendant de α , avec un δ_n bien choisi pour garantir que l'IC soit non-asymptotiquement conservateur et tel que $\lim_{n\to+\infty} \delta_n = 0$ pour avoir l'exactitude asymptotique est intéressante. C'est justement ce que permettent d'obtenir l'inégalité de Berry-Esseen et les expansions d'Edgeworth.
Inégalité de Berry-Esseen et expansions d'Edgeworth Ce paragraphe présente brièvement l'inégalité de Berry-Esseen (BE) et les expansions d'Edgeworth (EE).²² Pour commencer, il peut être intéressant de rappeler quelques résultats asymptotiques et non-asymptotiques pour situer l'inégalité de BE parmi ceux-ci.

La Loi (faible) des Grands Nombres (LGN) est un résultat asymptotique qui donne $\overline{D}_n \xrightarrow{P} \mathbb{E}[D]$. Les inégalités de concentration (dont les inégalités de BC et de H sont des exemples) sont des résultats non-asymptotiques : pour toute taille d'échantillon n, elles fournissent un nombre $t_{n,\alpha}$ tel que l'événement $\{|\overline{D}_n - \mathbb{E}[D]| \le t_{n,\alpha}\}$ arrive avec une grande probabilité (au moins $1 - \alpha$). Elles peuvent être vues comme une quantification, un équivalent non-asymptotique de la LGN.

Le Théorème Central Limite (TCL) donne une approximation de la distribution de \overline{D}_n lorsque la taille de l'échantillon tend vers l'infini : de façon informelle, on a $\overline{D}_n \sim \mathcal{N}(\mathbb{E}[D], \mathbb{V}(D)/n)$ asymptotiquement. En ce sens, le TCL précise la vitesse de convergence de la LGN, l'habituel $1/\sqrt{n}$.²³

L'inégalité de BE borne la distance uniforme (en norme infinie) entre la fonction de répartition (f.d.r) de \overline{D}_n , correctement centré et normalisé, et Φ , la f.d.r d'une loi Normale centrée réduite $\mathcal{N}(0, 1)$:

$$\sup_{x \in \mathbb{R}} \left| \mathbb{P}\left(\sqrt{n} \, \frac{\overline{D}_n - \mathbb{E}[D]}{\sqrt{\sigma^2}} \le x \right) - \Phi(x) \right| \le \frac{1.88 \, \mathbb{E}(|D - \mathbb{E}[D]]|^3)}{\sigma^3 \sqrt{n}},$$

où σ^2 est la variance de D.²⁴ Ainsi, l'inégalité de BE peut être vue comme une quantification, un équivalent non-asymptotique du TCL. Il faut noter que, par rapport au TCL, l'inégalité de BE requiert un moment supplémentaire : D doit admettre un moment d'ordre trois fini.

Introduisons quelques notations supplémentaires pour présenter les EE. On utilise ici les notations du chapitre 4. Soit $(X_i)_{i=1,...,n}$ des variables aléatoires réelles indépendantes mais non nécessairement identiquement distribuées (i.n.i.d.) supposées être centrées et ayant un moment d'ordre quatre fini : $\gamma_i := \mathbb{E}[X_i^4] < +\infty$. On pose

- l'écart-type de la somme des X_i , $B_n := \sqrt{\sum_{i=1}^n \mathbb{E}[X_i^2]}$, ainsi, la somme standardisée des X_i est $S_n := \sum_{i=1}^n X_i/B_n$,
- l'écart-type individuel moyen $\overline{B}_n := B_n / \sqrt{n}$,
- l'écart-type individuel $\sigma_i := \sqrt{\mathbb{E}[X_i^2]},$
- la moyenne des moments ordinaires standardisés d'ordre trois $\lambda_{3,n} := \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}[X_i^3] / \overline{B}_{n}^3$
- la moyenne des moments absolus standardisés d'ordre $p K_{p,n} := \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}[|X_i|^p]/(\overline{B}_n)^p$, pour $p \in \mathbb{N}^*$.

²²Le chapitre 4 donne plus de détails.

 $^{^{23}\}mathrm{On}$ peut remarquer que les inégalités de concentration indiquent également la vites se de convergence de la LGN.

 $^{^{24}}$ L'inégalité de BE écrite ici utilise la constante originale de 1.88 de Berry (1941). Des travaux plus récents ont obtenu de plus petites constantes. Une des contributions du chapitre 4 est d'améliorer cette inégalité en supposant des moments d'ordre quatre finis.

Dans le cas de données *i.i.d.*, on peut omettre l'indice *i* pour alléger les notations, et on a $\sigma_i = \sigma$ l'écart-type de X, $B_n = \sigma \sqrt{n}$, $K_{4,n} = K_4 = \mathbb{E}[|X|^4]/\sigma^4$ le kurtosis de X (une mesure de l'épaisseur des queues de distribution) et $S_n = \sqrt{n}(\overline{X}_n - \mathbb{E}[X])/\sigma$.

Les inégalités de Berry-Esseen cherchent à majorer la quantité

$$\Delta_{n,\mathrm{B}} := \sup_{x \in \mathbb{R}} \left| \mathbb{P} \left(S_n \le x \right) - \Phi(x) \right|$$

Les expansions d'Edgeworth sont un raffinement des inégalités de BE prenant en compte la présence d'asymétrie non-asymptotique dans la distribution de S_n . Ils cherchent à contrôler

$$\Delta_{n,\mathrm{E}} := \sup_{x \in \mathbb{R}} \left| \mathbb{P}\left(S_n \le x \right) - \Phi(x) - \underbrace{\frac{\lambda_{3,n}}{6\sqrt{n}} (1 - x^2)\phi(x)}_{=:\mathrm{Edg}_n(x)} \right|,$$

où $\Phi(x) + \frac{\lambda_{3,n}}{6\sqrt{n}}(1-x^2)\phi(x)$ est appelé l'expansion d'Edgeworth au premier ordre de $\mathbb{P}(S_n \leq x)$ (ϕ est la densité d'une $\mathcal{N}(0, 1)$, la dérivée de Φ).

On peut se demander pourquoi s'intéresser à des bornes sur les quantités $\Delta_{n,B}$ et $\Delta_{n,E}$. La réponse est qu'elles donnent les inégalités de concentration idoines pour obtenir des ICs non-asymptotiquement conservateurs mais simultanément asymptotiquement exacts. En effet, pour tout $n \in \mathbb{N}^*$, pour tout t > 0, on a (union disjointe et passage à l'événement complémentaire)

$$\begin{split} \mathbb{P}\left(\frac{\sqrt{n}\left|\overline{X}_{n}\right|}{\sigma} > t\right) &\leq 1 - \mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq t\right) + \mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq -t\right) \\ (+ \text{ et - pour voir un EE}) &\leq 1 - \left(\mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq t\right) - \Phi(t) - \text{Edg}_{n}(t)\right) \\ &+ \left(\mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq -t\right) - \Phi(-t) - \text{Edg}_{n}(-t)\right) \\ &- \Phi(t) - \text{Edg}_{n}(t) + \Phi(-t) + \text{Edg}_{n}(-t) \end{split}$$

(inég. triangulaire, $\text{Edg}_{n}(\cdot)$ paire, $\leq 2\Phi(-t) + \left|\mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq t\right) - \Phi(t) - \text{Edg}_{n}(t)\right| \\ &\text{ et } \Phi(t) = 1 - \Phi(-t)) + \left|\mathbb{P}\left(\frac{\sqrt{n}\,\overline{X}_{n}}{\sigma} \leq -t\right) - \Phi(-t) - \text{Edg}_{n}(-t)\right| \\ &\quad (\text{définition de } \Delta_{n,\text{E}}) \leq 2\left\{\Phi(-t) + \Delta_{n,\text{E}}\right\}. \end{split}$

Finalement, si on résout en t l'équation

$$2\left\{\Phi(-t) + \Delta_{n,E}\right\} = \alpha \iff \Phi(t) = 1 - \frac{\alpha}{2} + \Delta_{n,E},$$

et si on pose $X_i := D_i - \mathbb{E}[D]$, en prenant l'événement complémentaire dans le précédent calcul, on obtient

$$\Pr\left(\frac{\sqrt{n}\left|\overline{D}_{n}-\mathbb{E}[D]\right|}{\sigma} \le q_{\mathcal{N}(0,1)}\left(1-\frac{\alpha}{2}+\Delta_{n,\mathrm{E}}\right)\right) \ge 1-\alpha.$$
(8)

Si on suppose que σ , l'écart-type de X (ou de façon équivalente, de D) est connu (cas de l'inférence sur une espérance avec une variance supposée connue), dès lors qu'on dispose

d'une borne δ_n sur $\Delta_{n,E}$ qui est telle que $\lim_{n\to+\infty} \delta_n = 0$, ce calcul permet d'obtenir un CI ayant les propriétés recherchées : (i) le résultat de l'équation (8) est valide pour toute taille d'échantillon, d'où un IC non-asymptotiquement conservateur ; (ii) la longueur de l'IC converge vers celle de $\operatorname{CI}_{as}(1-\alpha, n)$, d'où un IC asymptotiquement exact.

Lorsque σ est inconnu, une étape supplémentaire est requise afin de le remplacer par son estimateur $\sqrt{\hat{\sigma}_n^2}$, mais l'idée générale demeure la même.

En substance, le chapitre 5 réalise la même analyse pour l'estimateur des Moindres Carrés Ordinaires (MCO) dans un modèle de régression linéaire au lieu de la moyenne empirique \overline{D}_n étudiée dans cet exemple introductif.

L'exactitude asymptotique est une propriété théorique intéressante. Toutefois, en pratique, pour un niveau nominal donné, plus la borne δ_n sur $\Delta_{n,E}$ est petite, plus la longueur de l'IC est petite, d'où une meilleure précision. De là l'intérêt de trouver des constantes aussi petites que possibles pour les bornes sur $\Delta_{n,B}$ et $\Delta_{n,E}$. A ma connaissance, les meilleures bornes existantes viennent de Shevtsova (2013), avec $\Delta_{n,B} \leq 0.5583K_{3,n}/\sqrt{n}$ (cas i.n.i.d.) et $\Delta_{n,B} \leq 0.4690K_{3,n}/\sqrt{n}$ (cas i.i.d.). Elles sont obtenues en supposant seulement des moments d'ordre trois finis.

Dans l'étude des régressions linéaires, la façon standard de faire de l'inférence, à partir de la normalité asymptotique de l'estimateur MCO, nécessite des moments d'ordre quatre finis pour les régresseurs. C'est pourquoi le chapitre 4 cherche à améliorer les bornes sous l'hypothèse de moments d'ordre quatre finis. Nous faisons cela en obtenant des bornes sur $\Delta_{n,E}$, qui donnent des bornes sur $\Delta_{n,B}$. En effet, l'inégalité triangulaire et sup_{$x \in \mathbb{R}$}($|1 - x^2|\phi(x)/6$) < 0.0665 permettent d'écrire

$$\Delta_{n,\mathrm{B}} \le \Delta_{n,\mathrm{E}} + \frac{0.0665\,\lambda_{3,n}}{\sqrt{n}}.$$

Finalement, il faut garder en tête que les intervalles de confiance $\operatorname{CI}_{BC}(1 - \alpha, n)$ (respectivement $\operatorname{CI}_H(1 - \alpha, n)$) nécessitent une borne M connue (des bornes a et b connues) sur la variance (le support) des observations. De même, l'IC obtenu ci-dessus à partir des inégalités de Berry-Esseen et des expansions d'Edgeworth requièrent des bornes sur certains moments afin de calculer δ_n , à savoir sur $\lambda_{3,n}$ (le coefficient d'asymétrie de X dans un cadre *i.i.d.*), $K_{3,n}$ (le moment absolu standardisé d'ordre trois) et $K_{4,n}$ (le kurtosis de X dans un cadre *i.i.d.*). Il s'avère qu'une borne sur $K_{4,n}$ est en fait suffisante puisque $|\lambda_{3,n}| < 0.621K_{3,n}$ (Pinelis, 2011) et $K_{3,n} \leq K_{4,n}^{3/4}$ (inégalité de Jensen). En l'absence d'hypothèses paramétriques fortes (comme l'hypothèse d'observations gaussiennes vue ci-dessus), il semble difficile pour l'inférence non-asymptotique d'éviter d'avoir à connaître de telles bornes. En pratique, il faut donc se demander comment les choisir. Les chapitres 4 et 5 discutent ce point.

Présentation et résumé de chacun des chapitres

Il s'agit des traductions en français des présentations et des résumés (abstract) présents dans le corps du manuscrit, rédigé en anglais, à chaque début de chapitre.

Chapitre 1 "segregsmall : une commande Stata pour estimer la ségrégation dans un contexte de petites unités".

Présentation. Ce premier chapitre est directement repris d'un article publié avec deux co-auteurs, Xavier D'Haultfœuille et Roland Rathelot dans *The Stata Journal*, 21(1), pages 152-179, mars 2021 (accès à l'article depuis le site internet de l'éditeur ici). Par rapport à la version publiée, la version de ce manuscrit présente simplement quelques modifications mineures et des annexes supplémentaires. Elle présente une commande Stata permettant d'estimer des indices de ségrégation. Le package Stata correspondant peut-être installé en exécutant dans une console Stata la ligne suivante :

net install segregsmall, from ("https://raw.githubusercontent.com/rolandrr/segregsmall-stata/master/")

Résumé. Considérons une population d'intérêt, formée d'un groupe minoritaire et d'un groupe majoritaire, dont les individus sont répartis dans des unités, telles que des quartiers résidentiels, des entreprises, des classes, etc. D'un point de vue qualitatif, il y a de la ségrégation dès lors que le processus d'allocation des individus dans les unités conduit à une concentration des individus minoritaires dans certaines unités davantage que dans d'autres. Les mesures quantitatives de la ségrégation sont confrontées au biais de petites unités ("small-unit bias"). En effet, quand les unités contiennent peu d'individus, les indices fondés sur les proportions observées d'individus minoritaires par unité sont biaisés. Ils sur-estiment le niveau réel de ségrégation, indiquant par exemple un niveau non nul de ségrégation alors même que le processus d'allocation serait entièrement aléatoire, répartissant les individus minoritaires uniformément entre les différentes unités. La commande Stata segregsmall met en oeuvre trois méthodes permettant de corriger ce biais : l'approche non-paramétrique, avec identification partielle, de D'Haultfœuille et Rathelot (Quantitative Economics, 2017), l'approche paramétrique de Rathelot (Journal of Business and Economic Statistics, 2012) et la correction linéaire de Carrington and Troske (Journal of Business and Economic Statistics, 1997). La commande permet également de réaliser des analyses conditionnelles, c'est-à-dire des mesures de la ségrégation prenant en compte les caractéristiques des individus ou des unités.

Chapitre 2 "Mesures de plusieurs dimensions de la ségrégation résidentielle en France entre 1968 et 2019 à partir des grappes de logements de l'Enquête Emploi".

Présentation. Le premier chapitre présente la commande Stata segregsmall. Ce second chapitre utilise cet outil pour mesurer la ségrégation résidentielle en France. L'analyse est faite entre 1968 et 2019 afin d'étudier l'évolution temporelle du niveau de ségrégation. Elle compare également les niveaux de ségrégations entre plusieurs dimensions, c'est-à-dire différentes façons (à partir de variables individuelles démographiques et socio-économiques) de définir les groupes minoritaire et majoritaire.

Je remercie l'INSEE, le réseau Quetelet-Progedo Diffusion et le Comité du Secret Statistique pour l'accès aux bases de données utilisées dans ce projet : Enquête Emploi 1968-2002, Enquête Emploi en continu (version FPR) 2003-2019.

Résumé. Cet article tire profit de la méthode de sondage de l'Enquête Emploi française (Labor Force Survey), qui échantillonne des grappes (clusters) d'une trentaine de logements

adjacents, pour étudier différentes dimensions de la ségrégation résidentielle en France entre 1968 et 2019. Ces grappes constituent des aires résidentielles ou voisinages pertinents pour étudier la ségrégation résidentielle sous réserve de prendre en compte le biais de petites unités ("small-unit bias") afin de pouvoir comparer de façon fiable les niveaux de ségrégation au cours du temps ou pour différentes dimensions (par exemple, personnes de nationalité française contre personnes de nationalité étrangère, personnes au chômage contre personnes employées sur le marché du travail, personnes ayant un diplôme universitaire contre les autres, cadres et professions intellectuelles supérieures contre les autres professions et catégories socio-professionnelles (PCS), etc.) En utilisant les méthodes développées par D'Haultfœuille et Rathelot (Quantitative Economics, 2017) et Rathelot (Journal of Business and Economic Statistics, 2012), l'article estime des indices de ségrégation annuels pour différentes définitions du groupe minoritaire et du groupe majoritaire, cherchant ainsi à quantifier et comparer différentes dimensions de la ségrégation résidentielle. Les résultats suggèrent deux conclusions principales. Premièrement, quelle que soit la dimension étudiée (ethnicité, immigrés contre non-immigrés, nationalités, PCS, positions sur le marché du travail, diplômes), les indices estimés ne présentent pas d'évolution significative au cours du temps : au sein de chaque dimension, le niveau de ségrégation résidentielle est ainsi resté globalement stable au cours des dernières décennies. Deuxièmement, les estimations mettent en évidence des niveaux de ségrégation différents selon la dimension étudiée, avec l'ordre suivant, par niveau décroissant de ségrégation : par nationalité, par ethnicité (en utilisant comme proxy le pays de naissance des parents d'un individu), entre immigrés et non-immigrés, puis par statuts sociaux (via les PCS ou les niveaux de diplôme), et enfin selon les positions sur le marché du travail (chômeurs contre employés). Une analyse conditionnelle, distinguant les grappes situées dans des aires urbaines de plus de 200 000 habitants de celles dans de plus petites aires urbaines, complète l'analyse inconditionnelle.

Chapitre 3 "Identification et estimation d'un indice de polarisation dans de grands ensembles de choix, avec une application aux débats du Congrès américain entre 1873 et 2016".

Présentation. Ce chapitre provient d'un projet en cours, un travail conjoint avec Xavier D'Haultfœuille et Roland Rathelot. Dans la lignée des chapitres 1 et 2, il traite de l'identification et de l'estimation de la polarisation. La différence est qu'il se concentre sur des situations de grands ensembles de choix (large or high-dimensional choice sets) au sens où il est possible d'observer certaines options, voire même une large proportion des options, qui ne sont choisies qu'une seule fois. Dans un contexte de ségrégation, cela correspondrait à des unités comportant un seul individu (une aire résidentielle formée d'une seule personne pour la ségrégation résidentielle, par exemple). De telles unités étaient très rares et exclues de l'analyse dans ces contextes de ségrégation (voir l'option withsingle de la commande Stata segregsmal1) tandis que cette possibilité est explicitement prise en compte et étudiée dans ce chapitre 3.

Résumé. Récemment, les divisions politiques ont semblé s'accroître dans plusieurs démocraties. Le langage, en tant que déterminant fondamental de l'identité d'un groupe, pourrait être un des facteurs de ces divisions. "Chasse aux sorcières" ("witch hunt") au lieu d'"audiences dans la procédure de destitution" ("impeachment hearing") ; "travailleurs sans-papiers" ("undocumented workers") contre "étrangers clandestins" ("illegal aliens") ; "charges sociales" à la place de "cotisations sociales" (voire "part du salaire mutualisé et différé") : ce genre d'expressions partisanes font référence aux mêmes objets mais

avec des connotations différentes. Elles se diffusent dans l'espace médiatique et peuvent influencer, par des effets de formulation notamment (framing effect), les opinions publiques. D'où l'intérêt de mesurer le niveau de polarisation du langage politique et de comparer son évolution temporelle ou entre pays. Une possibilité pour cela serait une étude linguistique des discours politiques. Une autre voie est l'étude statistique laquelle, malgré ses simplifications dans la représentation des textes (les données finales sont des décomptes de mots), permet d'étudier d'immenses corpus (exhaustifs) de textes sans avoir à spécifier ex-ante les expressions partisanes. Elle cherche à quantifier dans quelle mesure des groupes distincts (par exemple, les Démocrates et les Républicains au Congrès américain entre 1873 et 2016) utilisent des mots différents lorsqu'ils s'expriment. Gentzkow, Shapiro et Taddy (*Econometrica*, 2019) abordent ce problème au moyen d'un modèle de choix discret estimé grâce à une méthode d'apprentissage statistique de pénalisation. Nous proposons une méthode alternative présentant plusieurs avantages : (i) elle repose sur un résultat théorique d'identification partielle du paramètre d'intérêt (un indice de polarisation) dans un modèle statistique en partie testable, (ii) elle fournit des estimateurs simples et très peu coûteux numériquement des bornes d'identification ainsi qu'un intervalle de confiance pour le paramètre, (iii) elle requiert seulement des données agrégées (sans avoir à suivre l'identité des individus réalisant les choix). Par conséquent, il est facile d'appliquer notre méthode à d'autres contextes avec la même problématique de quantifier les différences entre les choix faits par des individus séparés en deux groupes dans un grand ensemble d'options possibles (high-dimensional or large choice sets) – c'est-à-dire lorsque le nombre d'options possibles est grand relativement au nombre de choix observés dans les données. Dans notre application, ces choix sont des choix de mots prononcés par des représentants Démocrates ou Républicains, mais il pourrait aussi bien s'agir de choix de localisation de résidence entre des immigrés et non-immigrés pour étudier une dimension de la ségrégation résidentielle par exemple, ou encore de produits choisis par deux groupes distincts de consommateurs.

Chapitre 4 "Bornes non-asymptotiques explicites sur la distance à l'expansion d'Edgeworth au premier ordre".

Présentation. Ce quatrième chapitre ouvre la seconde partie de cette thèse qui concerne l'inférence non-asymptotique. Il s'agit d'un travail commun avec deux co-auteurs, Alexis Derumigny et Yannick Guyonvarch. La version présentée dans cette thèse est une version révisée d'une version préprint Arxiv (arXiv:2101.05780v1).

Chronologiquement, ce travail pour borner des expansions d'Edgeworth (et ainsi obtenir des inégalités de type Berry-Esseen) débuta après une version initiale du chapitre 5, dans laquelle on utilisait les inégalités de Berry-Esseen existantes ne reposant que sur des moments d'ordre trois finis. Toutefois, dans les modèles économétriques de régressions linéaires, il est courant de supposer des moments d'ordre *quatre* finis pour les régresseurs afin d'avoir un estimateur consistant de la variance asymptotique de l'estimateur MCO, et d'obtenir ainsi les intervalles de confiance et tests statistiques asymptotiques usuels. Des bornes précises, avec des constantes numériques aussi petites que possibles, sont importantes en pratique pour faire de l'inférence non-asymptotique. D'où la motivation initiale de ce chapitre : améliorer les bornes de Berry-Esseen existantes lorsqu'on suppose des moments d'ordre quatre finis. Nous faisons cela au moyen de bornes sur les expansions d'Edgeworth, lesquelles permettent d'affiner les inégalités de Berry-Esseen en prenant en compte une possible asymétrie de la loi des données. Le projet s'est ensuite développé pour étudier les cadres *i.n.i.d.* et *i.i.d.* et propose également des bornes plus précises sous une hypothèse de régularité supplémentaire, qui, en substance, requiert d'avoir des observations dont la loi de probabilité admet une composante absolument continue par rapport à la mesure de Lebesgue (par opposition à des lois discrètes).

Résumé. Cette article étudie des bornes sur la distance uniforme (en norme infinie) entre la fonction de répartition d'une somme standardisée de variables aléatoires indépendantes et centrées admettant des moments d'ordre quatre finis et son expansion d'Edgeworth au premier ordre. Les bornes existantes sont améliorées dans deux cadres : lorsque les variables sont indépendantes mais non identiquement distribuées et lorsqu'elles sont indépendantes et identiquement distribuées. Des améliorations supplémentaires sont obtenues lorsque le moment d'ordre trois de la distribution est nul (cas de lois symétriques). L'article propose également des versions adaptées de ces bornes sous une hypothèse additionnelle de régularité qui concerne le comportement des queues de la fonction caractéristique de la somme normalisée des variables. Finalement, nous présentons une application de ces résultats qui explique l'absence de validité non-asymptotique (pour une taille d'échantillon finie) des tests unilatéraux reposant sur l'approximation gaussienne d'une moyenne empirique.

Chapitre 5 "Sur la construction d'intervalles de confiance non-asymptotiques dans les modèles linéaires".

Présentation. Ce cinquième chapitre utilise les résultats du quatrième chapitre pour construire des intervalles de confiance non-asymptotiques dans les modèles linéaires. Il s'agit à nouveau d'un travail commun avec Alexis Derumigny et Yannick Guyonvarch.

Résumé. Nous cherchons à construire des intervalles de confiance (ICs) nonasymptotiques pour les fonctionnelles linéaires du vecteur des coefficients d'un modèle de régression linéaire, c'est-à-dire, des ICs qui ont une probabilité de contenir la vraie valeur du paramètre au moins égale à leur niveau nominal quel que soit la taille de l'échantillon. Dans un cadre de régression linéaire avec des régresseurs exogènes, nous proposons un nouvel intervalle de confiance, ayant une expression explicite, et étroitement relié à l'IC standard reposant sur la t-statistique. Toutefois, contrairement à ce dernier dont la validité théorique n'est qu'asymptotique, notre IC est valide pour toute taille d'échantillon et cela uniformément sur une large classe de distributions caractérisée uniquement par des restrictions de moments (et non des hypothèses paramétriques). En particulier, nous autorisons l'hétéroscédasticité des résidus mais avons besoin d'une borne sur le kurtosis de la fonction d'influence de l'estimateur MCO. De plus, la longueur de notre IC converge en probabilité vers celle de l'IC standard reposant sur la t-statistique. Ainsi, notre IC est asymptotiquement exact. Nous étudions les performances pratiques de notre méthode par des simulations. Nous étendons également nos résultats théoriques aux modèles linéaires avec endogénéité. Dans le cadre d'un régresseur endogène, d'un instrument et de variables de contrôle exogènes, nous proposons une modification de la région de confiance d'Anderson-Rubin qui est valide non-asymptotiquement pour toute taille d'échantillon sous des conditions analogues au cas exogène. Notre travail s'appuie sur une littérature statistique ancienne, influencée de manière cruciale par Esseen (1945) et Cramer (1962).

Chapter 1

segregsmall: a Stata command to estimate segregation in the presence of small units

This first chapter is directly based on an article published with two co-authors, Xavier D'Haultfœuille and Roland Rathelot, in *The Stata Journal*, 21(1), pp. 152-179, March 2021 (online access to publisher's website here). The present chapter displays some minor modifications and additional appendices compared to the published version. It presents a Stata package to estimate segregation indices. The package can be installed from the Net by typing within Stata the following command:

net install segregsmall, from ("https://raw.githubusercontent.com/rolandrr/segregsmall-stata/master/")

Abstract Suppose that a population comprised of a minority and a majority group is allocated into units, such as neighborhoods, firms, school classes, etc. Qualitatively, there is some segregation whenever the allocation process leads to the concentration of minority individuals in some units more than in others. Quantitative measures of segregation have struggled with the small-unit bias. When units contain few individuals, indices based on the minority shares in units are upward biased. For instance, they would point to a positive amount of segregation even when the allocation process is strictly random. The Stata command **segregsmall** implements three recent methods correcting for such bias: the nonparametric, partial identification approach of D'Haultfœuille and Rathelot (*Quantitative Economics*, 2017), the parametric model of Rathelot (*Journal of Business and Economic Statistics*, 1997). The package also allows for conditional analyses, namely measures of segregation taking into account characteristics of the individuals or the units.

1.1 Introduction

We consider a population made of two groups (minority and majority) whose individuals are spread across units. Units can be geographical areas, residential neighborhoods, firms, school classes, or other clusters, provided that every individual belongs to exactly one unit. We seek to measure the extent to which individuals from the minority group are concentrated in some units more than in others. Throughout the paper, we follow the literature and use the word "segregation" as a neutral term to refer to such concentration. Measuring the magnitude of segregation is a necessary step to understand the underlying mechanisms and design adequate policies.

A natural way to measure segregation is to start from the minority shares X_i/K_i , where X_i is the number of individuals from the minority group and K_i the number of individuals (or unit's size) in unit $i \in \{1, ..., n\}$, and then compute an inequality index based on the distribution of the proportions X_i/K_i across the *n* units.

There are two possible benchmarks to assess the magnitude of these indices. Evenness relates to the case where all minority shares X_i/K_i are equal across units. Randomness relates to the case where the underlying allocation process assigns minority individuals at random across units. If p_i is the probability that an arbitrary individual in unit *i* belongs to the minority, randomness means that the probabilities p_i are equal across units. Past research has stressed the difference between both benchmarks, especially when the units are of small size (Cortese et al., 1976). The minority share X_i/K_i is only an estimate of p_i , and even if p_1, \ldots, p_n are all equal, there will be some variation in the X_i/K_i , all the more so as the units' sizes K_i are small. If one is interested in the deviations from the randomness case, indices based on minority shares, which measure the deviation from evenness, will overestimate the level of segregation. This issue is known as the small-unit bias.

The problem is pervasive in applied research. For workplace and school segregation, a large share of firms have less than ten employees, and classes usually have between twenty and forty students. The bias also arises when the units are not small *per se*, but only surveys of individuals are available. This is the case when one attempts to measure residential segregation using the local strata of households surveys.

Two main approaches have been proposed in the literature to deal with the smallunit bias. One strand proposes to correct the so-called naive inequality indices based on the minority shares X_i/K_i . The idea was initially proposed by Cortese et al. (1976) and Winship (1977) for the Duncan index. Carrington and Troske (1997, CT hereafter) extend the correction to other indices. Åslund and Skans (2009) adapt it to measure segregation conditional on covariates. Allen et al. (2015) develop another adjustment based on bootstrap. These corrections all aim to switch the benchmark from evenness to randomness by subtracting an estimate of the bias from the initial, naive index.

Another approach, adopted by Rathelot (2012, R hereafter) and D'Haultfœuille and Rathelot (2017, HR hereafter), defines segregation using an inequality index based on the unobserved probabilities p_i , as a functional of the distribution F_p of p_i . In line with the rest of the literature, they assume that the X_i are independent and follow a Bin (K_i, p_i) distribution. Conditional on K_i and p_i , R assumes a mixture of Beta distributions for F_p and derives the segregation index as a function of the distribution parameters. HR follow a nonparametric method leaving F_p unspecified; they show that the first moments of F_p are identified under the previous binomial assumption and obtain partial identification results on the segregation measure. Both R and HR construct confidence intervals for the segregation indices. HR also extend the methodology to study conditional segregation indices, namely measures of "net" or "residual" segregation taking into account other covariates (either of units or individuals) that may influence the allocation process.

The Stata command segregsmall allows social researchers to measure segregation in the context of small units. The command implements the methods proposed by R, HR, and CT. Conditional indices are available for all three methods. With R and HR, the command computes confidence intervals obtained by bootstrap. Finally, the command also implements a test of the binomial assumption.

This paper describes the command and presents the three methods it implements. Section 1.2 defines the set-up, the parameters of interest and synthesizes the estimation and inference methods of R, HR, and CT. Section 1.3 details the syntax, options, stored results of the segregsmall command, and discusses its execution time. Section 1.4 presents an application of the command on French firm data to measure workplace segregation between foreigners and natives across workplaces. Section 1.5 concludes.

1.2 Set-up, estimation, and inference

1.2.1 The setting and the parameters of interest

The population studied is assumed to be split into two groups: a group of interest, henceforth the minority group, and the rest of the population.¹ Individuals are distributed across units. For each unit, we assume that there exists a random variable p that represents the probability for any individual belonging to this unit to be a member of the minority. The total number of individuals in a unit is denoted by K.

We now introduce the segregation indices we focus on hereafter. We consider first unconditional indices; conditional indices are introduced in Section 1.2.6. Let us first assume that K is fixed. A segregation index θ is then a functional of the cumulative distribution function (c.d.f) F_p of p and of $m_{01} = E(p)$, that is $\theta = g(F_p, m_{01})$.² Roughly speaking, one expects such an index to be minimal when F_p is degenerate (Dirac), and maximal when $p \in \{0, 1\}$ (Bernoulli). In the former case, the probability of belonging to the minority is the same in all units. In contrast, the minority group is concentrated in a subset of units only in the latter case.

The command segregsmall estimates five classical segregation indices satisfying this property, namely

$$\theta^{\rm D} = \frac{\int |u - m_{01}| \, \mathrm{d}F_p(u)}{2m_{01}(1 - m_{01})}$$
(Duncan)
$$\theta^{\rm T} = 1 - \frac{\int \{u \ln(u) + (1 - u) \ln(1 - u)\} \, \mathrm{d}F_p(u)}{m_{01} \ln(m_{01}) + (1 - m_{01}) \ln(1 - m_{01})}$$
(Theil),

¹As most of the literature, we restrict to two groups. The study of segregation with more than two groups is appealing for some applications but raises questions regarding what we want to measure as segregation in multi-minority situations. Reardon and Firebaugh (2002) address this issue and propose several multi-group segregation indices. Some can be written as convex combinations of two-group segregation indices over the different groups included in the analysis. For such indices, two-group methods can thus be extended to multi-group settings.

²Such a notation may seem redundant since m_{01} already depends on F_p , the reason why we make the dependence on m_{01} explicit will become clearer below.

$$\begin{aligned} \theta^{A(b)} &= 1 - \frac{m_{01}^{\frac{--}{1-b}}}{1-m_{01}} \left\{ \int (1-u)^{1-b} u^{b} dF_{p}(u) \right\}^{\frac{1}{1-b}} & \text{(Atkinson, with } b \in (0,1)), \\ \theta^{CW} &= \frac{\int (u-m_{01})^{2} dF_{p}(u)}{m_{01}(1-m_{01})} & \text{(Coworker)}, \\ \theta^{G} &= \frac{1-m_{01} - \int F_{p}^{2}(u) du}{m_{01}(1-m_{01})} & \text{(Gini)}. \end{aligned}$$

When K is random and takes values in \mathcal{K} , θ is defined as a weighted average of indices conditional on K = k, denoted $\theta^k = g(F_p^k, m_{01}^k)$ with F_p^k the c.d.f of p conditional on K = k, and $m_{01}^k = \mathbb{E}[p | K = k]$; hence, θ depends on the joint distribution of p and K. Whether we study segregation at the unit-level or at the individual-level matters for the weights used. The unit-level index θ_u satisfies

$$\theta_u = \sum_{k \in \mathcal{K}} \Pr(K = k) \theta^k, \tag{1.1}$$

whereas the individual-level segregation index θ_i is defined by

$$\theta_i = \sum_{k \in \mathcal{K}} \frac{k \Pr(K = k)}{\mathcal{E}(K)} \theta^k.$$
(1.2)

The small-unit bias To estimate θ , we assume hereafter that the researcher has at her disposal K; however, the probability p remains unobserved. Instead, she only observes X, the number of individuals belonging to the minority in the unit. By definition of p, we have E[X | K, p] = Kp, which implies that the proportion of individuals from the minority, X/K, is an unbiased estimator of p. However, because it varies conditional on p, X/K is more dispersed than p. As a result, we have for usual segregation indices, including the five ones above,

$$g(F_{X/K}, m_{01}) > g(F_p, m_{01}) = \theta.$$

In other words, even in the absence of statistical uncertainty on the distribution of X/K, we would still overestimate the segregation index by using X/K in place of p. Moreover, this bias increases as K decreases. We refer to this issue as the small-unit bias hereafter.³

The binomial assumption We assume henceforth that individuals are allocated into units independently from each other. Namely, X is assumed to follow, conditional on p and K, a binomial distribution Bin(K, p). This hypothesis may be restrictive when the allocation process is in some way sequential and influenced by the composition of units. But importantly, this assumption is testable (see Section 1.2.5).

1.2.2 Nonparametric approach

Identification This approach, followed by HR, leaves the distribution F_p of p unrestricted. Combined with the binomial assumption, it entails a nonparametric binomial mixture model for X. Let us first suppose that K is constant; if not, we can simply retrieve

³Appendix 1.A reproduces the simulation studies of Cortese et al. (1976) and of CT. They underscore the importance of the bias in empirical settings. Appendix 1.B gives explicit expressions of the plug-in estimators of $g(F_{X/K}, m_{01})$ for the five classical indices we consider – the Duncan or dissimilarity index, the Theil index, the Atkinson index, the Coworker index, and the Gini index.

aggregated indices θ_u and θ_i using (1.1) and (1.2). We also assume that K > 1; if K = 1, the distribution of X is not informative on θ and we only get trivial bounds on it, namely 0 and 1 for the five indices above.

First, some algebra yields a one-to-one mapping between the distribution of X, defined by the K probabilities $P_0 = (P_{01}, \ldots, P_{0K})'$ with $P_{0j} = \Pr(X = j)$, and the first K moments of F_p , denoted $m_0 = (m_{01}, \ldots, m_{0k})'$:

$$P_0 = Qm_0,$$

with Q the $K \times K$ matrix with generic entry (i, j) equal to $\binom{K}{j}\binom{j}{i}(-1)^{j-i}$. It follows that m_0 is identified from the distribution of X; hence any parameter depending only on m_0 is point identified. It is the case of θ^{CW} as soon as $K \ge 2$. Second, there may be a single distribution F^* corresponding to m_0 . This happens if (and only if) m_0 belongs to the boundary $\partial \mathcal{M}$ of the moment space \mathcal{M}^4 . Then F^* is a discrete distribution with at most L+1 support points, where L is the integer part of (K+1)/2. For instance, when K = 2, $\mathcal{M} = \{(m_{01}, m_{02}) \in [0, 1]^2 : m_{01}^2 \le m_{02} \le m_{01}\}$, since $V(p) \ge 0$ and $p^2 \leq p$. Then $\partial \mathcal{M}$ corresponds to Dirac and Bernoulli distributions, for which we have respectively V(p) = 0 and $p^2 = p$.

When m_0 belongs to the interior \mathcal{M} of the moment space, there are infinitely many distributions F_p corresponding to m_0 . Then, unless we consider θ^{CW} , θ is not identified in general. Nevertheless, HR show that the sharp identified set on θ can be computed in a relatively easy way under the following restriction.

Assumption 1.1. $g(F, m_{01}) = \nu \left(\int h(x, m_{01}) dF(x), m_{01} \right)$, where h and ν are continuous and $\nu(\cdot, m_{01})$ is monotonic.

Assumption 1.1 fails for the Gini but is satisfied by the Duncan, the Theil, the Atkinson, and the Coworker indices. Under this condition, the bounds on $\int h(x, m_{01}) dF(x)$, and thus on θ , are attained by distributions with no more than K + 1 support points.⁵ Specifically, let $\mathcal{D}_{m_0}^{K+1}$ denote the set of distributions on [0, 1] with at most K+1 support points for which the vector of first K moments equals m_0 . Then the sharp identified set on θ is $[\underline{\theta}, \overline{\theta}]$, with

$$\underline{\theta} = \inf_{F \in \mathcal{D}_{m_0}^{K+1}} g\left(F, m_{01}\right), \quad \overline{\theta} = \sup_{F \in \mathcal{D}_{m_0}^{K+1}} g\left(F, m_{01}\right).$$
(1.3)

The fact that, under Assumption 1.1, the sharp bounds are attained by discrete distributions with a limited number of support points (at most K + 1) makes the optimization problem feasible in practice. Indeed, $\mathcal{D}_{m_0}^{K+1}$ can be seen as a subset of $[0,1]^{2K+1}$ since a discrete distribution is characterized by its support points (nodes) and their associated probabilities (masses).

The following theorem, which reproduces Theorem 2.1 of HR, summarizes the previous discussion. Hereafter, we let $\underline{\theta}$ and $\overline{\theta}$ denote the sharp lower and upper bounds on θ , whether or not θ is point identified.

Theorem 1.1. – If $m_0 \in \partial \mathcal{M}$, $\underline{\theta} = \overline{\theta} = g(F^*, m_{01})$, where F^* is the unique c.d.f for which the first K moments are equal to m_0 . Moreover, F^* has at most L + 1 support points. - If $m_0 \in \mathcal{M}$ and Assumption 1.1 holds, $\underline{\theta}$ and $\overline{\theta}$ are defined by (1.3).

⁴This claim and several others of this section are proved in Krein and Nudel'man (1977). For that specific point and the following, see notably Theorem III.4.1 (another reference is Theorem I.2.5 in Dette and Studden (1997)).

⁵See HR and Theorem I.3.6 in Krein and Nudel'man (1977) for more details.

In the interior case, computing the bounds still requires a nonlinear optimization under constraints that are also nonlinear in the support points. Yet, the problem can be further simplified under additional assumptions using the theory of Chebyshev systems. In particular, it requires that the function h in Assumption 1.1 does not depend on m_{01} , a condition satisfied by the Theil and Atkinson indices. Basically, no numerical optimization is needed for those two indices to compute the bounds $\underline{\theta}$ and $\overline{\theta}$. The idea behind this is that the bounds are attained by two special discrete distributions, called *principal representations*. The interest is that finding the principal representations boils down to obtaining the roots of specific polynomials, which is much simpler and faster than solving (1.3). We refer to HR for more details on that matter.

Estimation Let us assume to have in hand an independent and identically distributed (i.i.d.) sample $(X_i)_{i=1,...,n}$ of n units, with constant sizes equal to K > 1. Theorem 1.1 shows that θ is either point or partially identified, depending on whether $m_0 \in \partial \mathcal{M}$ or $m_0 \in \overset{\circ}{\mathcal{M}}$. We follow this result to estimate the identified set $(\underline{\theta}, \overline{\theta})$. In a first step, we estimate P_0 , and thus $m_0 = Q^{-1}P_0$, by constrained maximum likelihood. The constraints come from the binomial mixture model: $P_0 \in \mathcal{P} = \{Qm : m \in \mathcal{M}\}$. To compute the constrained MLE, HR show Lemma 1.1 below. To state the result, let us define $N_k =$ $\sum_{i=1}^n \mathbbm{1}\{X_i = k\}$, the number of units in the sample with k minority individuals, the set of nodes $\mathcal{S}_{L+1} = \{(x_1, \ldots, x_{L+1}) : 0 \leq x_1 < \ldots < x_{L+1} \leq 1\}$ and the associated masses $\mathcal{T}_{L+1} = \{(y_1, \ldots, y_{L+1}) \in [0, 1]^{L+1} : \sum_{k=1}^{L+1} y_k = 1\}$.

Lemma 1.1. The constrained MLE $\hat{P} = (\hat{P}_1, \ldots, \hat{P}_K)'$ satisfies

$$\widehat{P}_{k} = \binom{K}{k} \sum_{j=1}^{L+1} \widehat{y}_{j} \widehat{x}_{j}^{k} (1 - \widehat{x}_{j})^{K-k}, \quad \forall k \in \{1, \dots, K\},$$

where $\hat{x} = (\hat{x}_1, \dots, \hat{x}_{L+1})$ and $\hat{y} = (\hat{y}_1, \dots, \hat{y}_{L+1})$ are given by

$$(\hat{x}, \hat{y}) = \operatorname*{argmax}_{(x,y)\in\mathcal{S}_{L+1}\times\mathcal{T}_{L+1}} \sum_{k=0}^{K} N_k \ln\left\{\sum_{j=1}^{L+1} y_j x_j^k (1-x_j)^{K-k}\right\}$$

In a second step, we estimate $(\underline{\theta}, \overline{\theta})$. First, we estimate m_0 by $\widehat{m} = Q^{-1}\widehat{P}$, where \widehat{P} is obtained in practice thanks to Lemma 1.1. Then, we check whether the estimator $\widehat{m} \in \partial \mathcal{M}$. A simple possibility to do so is checking whether the unconstrained MLE $\widetilde{P} = (\widetilde{P}_1, \ldots, \widetilde{P}_K)'$ satisfies $\widetilde{P} = \widehat{P}$ (in which case $\widehat{m} \in \mathcal{M}$ with probability approaching one) or not.⁶ Note that the unconstrained MLE simply satisfies $\widetilde{P}_k = N_k/n$ for all k.

When $\tilde{P} \neq \hat{P}$, we simply let $\underline{\hat{\theta}} = \overline{\hat{\theta}} = g(\hat{F}, \widehat{m}_1)$, where \hat{F} is the distribution corresponding to (\hat{x}, \hat{y}) . We refer to this situation as the *constrained case*.

If $\hat{P} = \hat{P}$, there are infinitely many distributions corresponding to \hat{m} and we estimate bounds for θ . We refer to this situation as the *unconstrained case*. For the Theil and Atkinson indices, the estimated bounds are obtained from the principal representations

⁶The intuition is the following. When $\tilde{P} \notin \mathcal{P}$, the constrained estimator \hat{m} is forced to be in \mathcal{M} and will end up in the boundary $\partial \mathcal{M}$. On the contrary, if $\tilde{P} \in \mathcal{P}$, $\tilde{P} = \hat{P}$, hence the constrained estimator \hat{m} and the unconstrained one $\tilde{m} = Q^{-1}\tilde{P}$ coincide and belongs to \mathcal{M} with probability approaching one when the sample size n goes to infinity. The procedure amounts to test whether a given vector of moment belongs to the moment space, a problem that has been studied extensively (the formal test is defined in Appendix D.1 of HR).

computed from \widehat{m} . For the Duncan index, optimization is required to obtain the estimated bounds. We obtain estimators of $\underline{\theta}$ and $\overline{\theta}$ by solving the optimization problems (1.3), replacing m_0 by its estimator \widehat{m} . Finally, the Coworker index only depends on (m_{01}, m_{02}) . Thus, whether or not $\widetilde{P} = \widehat{P}$, this index can be estimated directly by replacing (m_{01}, m_{02}) by $(\widehat{m}_1, \widehat{m}_2)$.

Inference When Assumption 1.1 holds, HR show that the estimators of the bounds are consistent: $(\hat{\theta}, \hat{\theta}) \xrightarrow{P} (\hat{\theta}, \overline{\theta})$ as the number of units *n* tends to infinity. Under additional assumptions, HR characterize their asymptotic distributions. This enables to build valid asymptotic confidence intervals (CIs) for the index θ using a modified bootstrap procedure. The construction needs to take into account the fact that the lower bound and upper bound collapse when $m_0 \in \partial \mathcal{M}$ (point-identification) whereas they differ when $m_0 \in \overset{\circ}{\mathcal{M}}$ (partial identification). The underlying idea relates to the construction of CIs in the case of partial identification (see Imbens and Manski (2004), Stoye (2009)). HR define a confidence interval for the interior case, where only one of the two ends of the interval matters in the asymptotic coverage in all situations, HR define the final confidence interval by selecting one of them according to the length of the estimated identification interval $(\hat{\theta} - \hat{\theta})$ relative to sampling error.⁷

Random unit size The previous identification and estimation results can be adapted to cases where K is random and takes values in a set \mathcal{K} . Using the definitions of θ_u and θ_i in (1.1) and (1.2), the idea is to reason conditional on the unit size to get each θ^k , $k \in \mathcal{K}$, and replace the theoretical weights by plug-in estimators. More precisely, let $\hat{\theta}^k$ and $\hat{\theta}^k$ denote the estimators of the bounds of θ^k based on the subsample of units of size k. Let $\widehat{\Pr}(K = k) = n^{-1} \sum_{i=1}^n \mathbb{1}\{K_i = k\}$ and $\widehat{\mathbb{E}(K)} = n^{-1} \sum_{i=1}^n K_i$. Then the estimators of the bounds on θ_u and θ_i satisfy

$$\begin{split} & \widehat{\underline{\theta}}_u = \sum_{k \in \mathcal{K}} \widehat{\Pr}(K = k) \underline{\widehat{\theta}}^k, \quad \widehat{\overline{\theta}}_u = \sum_{k \in \mathcal{K}} \widehat{\Pr}(K = k) \widehat{\overline{\theta}}^k, \\ & \underline{\widehat{\theta}}_i = \sum_{k \in \mathcal{K}} \frac{k \widehat{\Pr}(K = k)}{\widehat{\mathrm{E}(K)}} \underline{\widehat{\theta}}^k, \quad \widehat{\overline{\theta}}_i = \sum_{k \in \mathcal{K}} \frac{k \widehat{\Pr}(K = k)}{\widehat{\mathrm{E}(K)}} \overline{\widehat{\theta}}^k. \end{split}$$

Remark that as soon as for one unit size k the index θ^k is not point identified, the resulting aggregated index will be partially identified too. In other words, point identification of θ_u or θ_i requires to be in the constrained case for each $k \in \mathcal{K}$. This is unlikely to happen when the support of K contains very small sizes k, typically lower than 10.

Similar to the constant unit case, confidence intervals for the aggregated indices θ_u and θ_i are constructed by the modified bootstrap procedure detailed in HR. The randomness of K just involves an additional step that consists in drawing K in its empirical distribution.

Assuming independence between K and p The previous estimation and inference procedures are fully agnostic as regards possible dependence between K and p, which is a

⁷Essentially, when this length is large (resp. small) relative to sampling error, the uncertainty related to partial identification (resp. to sampling) prevails and the interior-type (resp. boundary-type) confidence interval is used.

safe option when unit size may be a potential determinant of segregation. However, if one is ready to impose independence between these two variables, the identified bounds on $\theta_u = \theta_i$ get closer to each other. This is because the F_p^k coincide with the unconditional distribution of p. Thus, we can gather all units and identify the first \overline{K} moments of F_p , with $\overline{K} = \max(\mathcal{K})$. Estimation and inference are performed as in the case of constant unit size, with K replaced by \overline{K} . Thus, assuming independence between K and p improves identification since we identify more moments of F_p . It also leads to more accurate estimators since one estimates a single vector P on the whole sample, instead of doing so on each subsample $\{i : K_i = k\}$, for all $k \in \mathcal{K}$.

An important particular case occurs when only some individuals in the unit are observed (e.g., survey data). Imagine units are of size $(K_i)_{i=1,...,n}$ but that, for each unit *i*, only $n_{K,i}$ individuals are sampled and observed. We let X_i denotes the number of individuals belonging to the reference group in this subgroup of $n_{K,i}$ people. As previously, X_i follows a binomial distribution $Bin(n_{K,i}, p_i)$ conditional on p_i and $n_{K,i}$. The previous results apply by simply replacing the unit size K with the number n_K of individuals observed in each unit. Moreover, in such settings, it is usually plausible to assume that the random variable n_K is independent of p conditional on the unit size as n_K depends on the survey process, which, a priori, is orthogonal to the segregation phenomenon.

1.2.3 Parametric approach

This approach, followed by R, is similar to that of HR, except that it imposes a parametric restriction on F_p . Specifically, it is supposed to be a mixture of Beta distributions. Combined with the binomial assumption for the conditional distribution of X, the model becomes fully parametric and thus can be estimated by maximum likelihood. Therefore, the indices are point identified, contrary to the nonparametric approach of HR.

A concern might be that the parametric restriction leads to invalid results when the model is misspecified. However, R shows through simulations that segregation indices associated with various distributions, both continuous and discrete, are accurately proxied by his parametric approach.

Estimation and inference As in HR, we first assume that K is constant. Let $B(\cdot, \cdot)$ denote the beta function, c the number of components of the beta mixture, and $v = (\alpha_j, \beta_j, \lambda_j)_{j=1,...,c}$ the vector of parameters with $(\alpha_j, \beta_j) \in \mathbb{R}^*_+ \times \mathbb{R}^*_+$ the two shape parameters of the *j*-th Beta distribution and $\lambda_j \in [0, 1]$ its weight $(\sum_{j=1}^c \lambda_j = 1)$. The probability density function of p distributed as a c-component mixture of Beta distributions with parameters v is

$$f_{v}(t) = \sum_{j=1}^{c} \lambda_{j} \frac{t^{\alpha_{j}-1} (1-t)^{\beta_{j}-1}}{B(\alpha_{j}, \beta_{j})}, \quad \forall t \in [0, 1].$$

In this model, the probability that k individuals belong to the minority group can be written, after some algebra, as:

$$\Pr_{v}\left(X=k\right) = \binom{K}{k} \sum_{j=1}^{c} \lambda_{j} \frac{B(\alpha_{j}+k,\beta_{j}+K-k)}{B(\alpha_{j},\beta_{j})}.$$

Thus, the log-likelihood satisfies, up to terms independent of the parameter v,

$$\ell(v) = \sum_{k=0}^{K} N_k \times \ln \left\{ \sum_{j=1}^{c} \lambda_j \frac{B(\alpha_j + k, \beta_j + K - k)}{B(\alpha_j, \beta_j)} \right\},\,$$

Maximizing $v \mapsto \ell(v)$ yields the maximum likelihood estimator \hat{v} . Using the parametric assumption on F_p , \hat{v} translates into an estimator \hat{F}_p of the distribution of p, which in turn yields an estimator $\hat{\theta}$ of θ . The explicit expressions of the five indices above, as functions of the parameter v, are given in Appendix 1.C. Inference can be achieved by the delta method or by the bootstrap, performed at the unit level.

Random unit size The adaptation to this case is exactly similar to HR method. For each $k \in \mathcal{K}$, the MLE of θ^k is obtained using the subsample of units of size k. The weights are estimated by their empirical counterparts. The estimated aggregated indices are then obtained by plug-in, using (1.1) and (1.2). When K and p are assumed independent, all units can be pooled, independently of their size, to compute the MLE of v for the whole sample. As above, the resulting estimator \hat{v} allows us to estimate the distribution of p, and then θ .

1.2.4 Correction of the naive index

The approaches of HR and R are immune to the small-unit bias as they directly estimate $g(F_p, m_{01})$. Other, previous approaches instead start from the naive index $\theta_N = g(F_{X/K}, m_{01})$ and attempt to modify it so that the parameter becomes less sensitive to changes in K. We present here the correction proposed by CT, which is the most popular in applied work.

CT's correction relies on the distinction between the randomness and evenness benchmarks, introduced notably by Cortese et al. (1976) and Winship (1977). Evenness corresponds to X/K being constant, whereas randomness refers to the case where p is constant. Under the binomial model, however, evenness cannot occur. The central idea of CT is then to convert θ_N , which measures departure from evenness, into a distance to randomness. To do so, CT compare θ_N to its expected value θ_N^{ra} under the random allocation of individuals into units.

Formally, let $\theta_N^{\rm ra}$ denote $g(F_{X^{\rm ra}/K}, m_{01})$, where $X^{\rm ra} | K \sim {\rm Bin}(K, E(X/K))$. $X^{\rm ra}/K$ is the proportion we would observe if p was constant and equal to E(p) = E(X/K). Then, assuming that $\theta \in [0, 1]$, a constraint satisfied by the five indices above, CT's correction θ_{CT} is defined by $\theta_{CT} = (\theta_N - \theta_N^{\rm ra})/(1 - \theta_N^{\rm ra})$. CT suggest the following simulation-based estimator of θ_{CT} . Let $\hat{E}(p)$ denote the sample average of X/K. For s = 1, ..., S, draw $X_{i,s}^{\rm ra} \sim {\rm Bin}(K_i, \hat{E}(p))$ independently for each unit i. Then, letting $\hat{F}_s^{\rm ra}$ and $\hat{m}_{1,s}$ denote respectively the empirical distribution and mean of $(X_{i,s}^{\rm ra}/K_i)_{i=1,...,n}$, compute $\hat{\theta}_{N,s}^{\rm ra} = g(\hat{F}_s^{\rm ra}, \hat{m}_{1,s})$. The estimator of $\theta_N^{\rm ra}$ is then the mean over the S replications, $\hat{\theta}_N^{\rm ra} = S^{-1} \sum_{s=1}^S \hat{\theta}_{N,s}^{\rm ra}$. Finally, $\hat{\theta}_{\rm CT} = (\hat{\theta}_N - \hat{\theta}_N^{\rm ra})/(1 - \hat{\theta}_N^{\rm ra})$, with $\hat{\theta}_N$ the plug-in estimator of θ_N . The quantiles of $(\hat{\theta}_{N,s}^{\rm ra})_{s=1,...,S}$ can be used to test that the data are consistent with random allocation using randomization tests (see Boisso et al., 1994, and CT).

Links with HR and R In general, $\theta_{CT} \neq \theta$. However, they do coincide in extreme cases of no segregation, where p is constant, and "full" segregation, where p follows a Bernoulli distribution. We refer to Section 2.3 of R and Section 2.4 of HR for further discussion on the relationship between θ_{CT} and θ .

1.2.5 Test of the binomial assumption

We have relied so far on the binomial assumption $X | K, p \sim Bin(K, p)$. This assumption implies that $P_0 \in \mathcal{P} = \{Qm : m \in \mathcal{M}\}$. A vector $(m_1, ..., m_K)$ in \mathcal{M} has to satisfy some restrictions, such as $m_2 \geq m_1^2$ (i.e., non-negative variance). Hence, we could have $Q^{-1}P_0 \notin \mathcal{M}$ if the distribution of X conditional on K and p is not binomial. In other words, the binomial assumption is testable.

HR propose a likelihood ratio test of $P_0 \in \mathcal{P}$, where the constrained estimator under the null hypothesis is \hat{P} , whereas the unconstrained MLE is \tilde{P} . Note that these estimators are already computed to estimate $(\underline{\theta}, \overline{\theta})$. For a unit size equal to k, the test statistic satisfies

$$LR_k = 2\sum_{x=0}^k N_x \ln\left(\frac{\tilde{P}_x}{\hat{P}_x}\right) = 2\sum_{x=0}^k N_x \ln\left(\frac{N_x}{n\hat{P}_x}\right),$$

where we let $N_x \ln[N_x/(n\hat{P}_x)] = 0$ if $N_x = 0$.

With a random unit size, the test statistic is then $LR_n = \sum_{k \in \mathcal{K}} \widehat{\Pr}(K = k) LR_k$, where in LR_k , $N_x = \sum_{i=1}^n \mathbb{1}\{K_i = k, X_i = x\}$. The critical values of the test are obtained by approximating the distribution of LR under the null by bootstrap. The bootstrap is performed as follows. First, we draw n units of sizes K_i^* in the empirical distribution of K. Second, we draw X_i^* according to $\widehat{P}^{K_i^*}$, where \widehat{P}^k is the constrained MLE of P_0^k , the distribution of X conditional on K = k. The bootstrapped test statistic LR^* is then computed in the sample $(K_i^*, X_i^*)_{i=1,\dots,n}$, which is drawn under the null hypothesis. For a level $1 - \alpha \in (0, 1)$, the critical region of the test is defined by:

$$CR = \left\{ LR > c_{1-\alpha} \left(LR^* \right) \right\},\,$$

with $c_{1-\alpha}(LR^*)$ the quantile of order $1-\alpha$ of LR^* .

The results of HR imply that the test has an asymptotic level equal to α and is consistent. Remark, however, that it tests $P_0 \in \mathcal{P}$, which is an implication of the binomial assumption, rather than this assumption itself. This means that the binomial assumption may fail but still, $P_0 \in \mathcal{P}$, that is, $X \mid K, p$ could fail to be binomial, yet the distribution of X given K could be rationalized by a binomial mixture.

1.2.6 Conditional segregation indices

Conditional indices aim at accounting for the fact that part of the segregation along the minority/majority dimension at stake may be driven by sorting according to other dimensions. In this sense, they measure the net or residual level of segregation when the contribution of covariates to segregation is removed (see Åslund and Skans, 2009). To illustrate this point, let us consider workplace segregation between foreigners and natives. Foreigners may be hired more in some sectors of the economy on the basis of sector-specific skills. Imagine an extreme case where, within each sector, all firms hire foreigners with the same probability. As long as these probabilities differ from one sector to another, an unconditional segregation index would be positive. On the contrary, the conditional index defined in (1.4) below would indicate no segregation as it controls for the influence of the sector, a characteristic of units, in the allocation process. Similarly, foreigners may be hired with the same probability for all low-skilled jobs (resp. all high-skilled jobs), but the probabilities for these two job types may differ. In this case, again, failing to account for this characteristic would lead to a positive unconditional index, while the conditional index defined in (1.5) below would indicate no segregation. The previous discussion underscores that covariates can be defined either at the unit level or at the level of an individual/position. We separate the two cases below, as they lead to different treatments.

Unit-level covariates Let $Z \in \{1, \ldots, \overline{Z}\}$ denote a characteristic of a unit, which is assumed to be discrete. To take into account Z in the allocation process, we measure segregation conditional on Z. For each $z \in \{1, \ldots, \overline{Z}\}$, let θ_{0z} denote the segregation index we consider conditional on Z = z. The subscript 0 indicates that we consider a generic index of interest, which could correspond to either θ (θ_u of θ_i with random K) or θ_{CT} . Whatever the index, the estimation of θ_{0z} is done exactly as in the unconditional case, focusing on the subsample $\{i : Z_i = z\}$.

The index θ_{0z} can be of interest by itself. We can also consider an aggregate conditional index defined as follows:⁸

$$\theta_{0,u}^{\text{cond}} = \sum_{z=1}^{\overline{Z}} \Pr(Z=z) \theta_{0z}.$$
(1.4)

The estimation of $\theta_{0,u}^{\text{cond}}$ is obtained by plug-in, with $n^{-1} \sum_{i=1}^{n} \mathbb{1}\{Z_i = z\}$ the empirical counterpart of $\Pr(Z = z)$. For HR and R methods, a similar bootstrap procedure as in the random size case provides asymptotic confidence intervals for $\theta_{0,u}^{\text{cond}}$.⁹

Individual- or position-level covariates Let $W \in \{1, \ldots, \overline{W}\}$ denote a characteristic of an individual or of a position. To resume the example of workplace segregation, a characteristic attached to individuals can be education, whereas a characteristic linked to positions can refer to the type of occupation (e.g., high-skilled versus low-skilled). While these two forms of covariates may lead to different interpretations, they are similar regarding estimation and inference.

For each unit and each type $w \in \{1, \ldots, \overline{W}\}$, we suppose to observe X_w and K_w , which are respectively the number of individuals with characteristic W = w (or in positions satisfying W = w) who belong to the minority group, and the overall number of individuals (or positions) of type W = w in the unit. As above, we define θ_{0w} as the segregation index of interest conditional on W = w. With individual- or position-level covariates, the idea is to consider the subsample of individuals (or positions) such that W = w, instead of a subsample of units. Hence, θ_{0w} can be estimated exactly as in the unconditional case simply using (X_w, K_w) instead of (X, K).¹⁰

Again, θ_{0w} might be a relevant parameter of interest on his own. Researchers can also be interested in an aggregated conditional index:

$$\theta_{0,i}^{\text{cond}} = \sum_{w=1}^{\overline{W}} \Pr(W = w) \theta_{0w}.$$
(1.5)

⁸Note that for θ_{CT} , the aggregate conditional indices defined by (1.4), and similarly for (1.5) below, slightly differ from the conditional index of Åslund and Skans (2009). Broadly speaking, (1.4) and (1.5) aggregate the corrected indices computed conditional on each type while Åslund and Skans (2009) do one unique correction in order to directly obtain their conditional corrected index. The former has the advantage of being more general and notably can be used as such in HR and R approaches.

⁹The initial step of the bootstrap procedure becomes drawing units in the joint empirical distribution of (K, Z).

¹⁰Remark that, in the general random sizes case without assuming $K \perp p$, it makes more sense to consider the index θ_i that uses individual-level weights (compared to unit-level ones) because the types are defined at this individual-/position-level.

The estimation of $\theta_{0,i}^{\text{cond}}$ is obtained by plug-in, with $(\sum_{i=1}^{n} K_{wi})/(\sum_{i=1}^{n} K_i)$ the empirical counterpart of $\Pr(W = w)$. For HR and R methods, as previously, a modified bootstrap procedure provides asymptotic confidence intervals for $\theta_{0,i}^{\text{cond}}$.¹¹

Finally, remark that conditional analyses with individual- or position-level covariates compound the small-unit bias issue, especially when the number \overline{W} of possible types is large. Indeed, the equivalent of a unit in unconditional analyses is now a cell unit \times type with X_w minority individuals among K_w individuals and, mechanically, $K_w \leq K$.

1.3 The segregsmall command

The segregsmall command is compatible with Stata 14.2 and later versions.

1.3.1 Syntax

The syntax of segregsmall is as follows:

```
segregsmall varlist [if] [in], method(string) format(string) [ conditional(string)
withsingle excludingsinglepertype independencekp level(#) repbootstrap(#) noci
testbinomial repct(#) atkinson(#)]
```

1.3.2 Description and main options

The command segregsmall estimates the five classical segregation indices mentioned above (Duncan, Theil, Atkinson, Coworker, and Gini) using D'Haultfœuille and Rathelot (2017), Rathelot (2012), or Carrington and Troske (1997) method. It provides confidence intervals obtained by bootstrap in the approaches of HR and R and allows for conditional analysis for all three methods.

- method specifies the method used. Its argument must be one of: np, beta, ct. Argument np, standing for nonparametric, implements HR method. The command does not report the Gini index in this case as it does not verify Assumption 1.1. The choice beta implements R's method assuming a Beta distribution for F_p .¹² Both methods provide estimates of the same parameters of interest, namely θ if K is fixed and, unless independencekp is specified, (θ_u, θ_i) if K is random. By default, they report asymptotic confidence intervals obtained by bootstrap. With the argument ct, the command estimates the naive and CT-corrected indices θ_N and θ_{CT} . Confidence intervals are not computed for these parameters.
- format indicates the format of the dataset used and needs to be either unit (datasets where an observation is a unit) or indiv (datasets where an observation is an individual). The option determines the variables to be put in varlist. For unconditional analyses (the default without option conditional), these are:
 - K X for unit-level datasets, K and X correspond to the variables K and X introduced in Section 1.2: the number of individuals and the number of minority individuals. K

¹¹The initial step of the bootstrap procedure becomes drawing units in the empirical distribution of units, hence keeping fixed the composition of the units with respect to W.

¹²R assumes a mixture of Beta distributions. However, simulations reveal that the differences between the indices obtained with a two or higher component mixture versus a simple Beta are marginal in most cases, **segregsmall** uses a Beta assumption for simplicity. Also, the command allows assessing the reliability of this restriction since the indices obtained with the beta restriction can be compared with the nonparametric estimates that leave F_p unrestricted.

has to be positive integers and $\tt X$ non-negative integers. Also, $\tt X$ should be lower or equal to $\tt K$ for each unit.^13

- id_unit I_minority for individual-level datasets, id_unit is the identifier of the unit the individual belongs to. I_minority is a dummy variable equal to 1 when the individual belongs to the minority group, 0 otherwise.
- conditional this option triggers the computation of conditional segregation indices. Its arguments must be either *unit* or *indiv* and it specifies the level at which are defined the covariates included in the analysis. For conditional analysis, *varlist* has to be:
 - K X Z for unit-level datasets, or id_unit I_minority Z for individual-level datasets, with covariates defined at unit-level (*unit*). The variables K, X, id_unit, and I_minority are the same as in unconditional analyses. Z corresponds to the variable Z, the characteristics of units defined in Section 1.2.6. Z needs to take values in $\{1, 2, \ldots, \overline{Z}\}$ with $\overline{Z} \ge 2$.
 - id_unit I_minority W for individual-level datasets with covariates defined at the level of individuals or any sub-unit level (indiv).¹⁴ W corresponds to the variable W, the individual (or position) characteristics introduced in Section 1.2.6. W has to take values in $\{1, 2, \ldots, \overline{W}\}$ with $\overline{W} \geq 2$.

1.3.3 Additional options

- withsingle includes single units (with only one individual) in the analysis. As explained in Section 1.2.2, single units are in general uninformative about the level of segregation. By default, they are not included in the data used. The option is available both for unconditional or conditional analyses.
- excludingsinglepertype excludes single cells (unit \times type) from the analysis. The option is only relevant and available in conditional analyses with covariates defined at the individual/position level. In this setting, the role of a unit in unconditional analyses is played by a cell defined as the intersection of a unit and an individual type (see Section 1.2.6). As just described, *units* with only one individual are dropped by default. Yet, this does not prevent the existence of single *cells* coming from units with more than one individual but that have only one individual with a given characteristic W = w. Without option excludingsinglepertype, those single cells are included in the analysis, which can lead to wide estimated identified intervals in HR method, all the more so as the number \overline{W} of types is large. With the option, they are dropped. For consistency, the options withsingle and excludingsinglepertype are mutually exclusive.
- independencekp assumes independence between K and p. The option is only available with np and *beta* methods. By default, the command is agnostic about potential links between K and p.
- level sets the confidence level, which has to be a scalar in (0, 1). With np and beta methods, by default, the traditional 90%, 95%, and 99% confidence levels are saved

¹³Besides, segregsmall requires all variables in *varlist* to be free of missing values. Missing values are not handled automatically to avoid possible confusions between zero and missing values for the variable X notably.

¹⁴Conditional analyses with individual- or position-level covariates logically require databases also defined at the individual-level.

(see Section 1.3.4) and the 95% confidence interval is displayed in Stata output. The option permits to save and display a personalized level besides (the other three are still stored). With ct method, by default, the empirical quantiles of the index under random allocation are stored for the orders 0.01, 0.05, 0.10, 0.90, 0.95, and 0.99. The option additionally saves the empirical quantiles at order τ and $1 - \tau$ with τ the argument of the option. Such quantities can be used to test the null hypothesis that the data are consistent with random allocation through randomization tests.

- repbootstrap specifies the number of bootstrap iterations used to construct confidence intervals in np and beta methods. The default number is 200. It is also the number of bootstrap repetitions used to test the binomial assumption.
- **noci** restricts the command to estimation: confidence intervals are not computed. The option is only applicable to *np* and *beta* methods.
- testbinomial implements the test of the binomial assumption. More precisely, with method(np) and without options independencekp nor noci, the test is made by default and saved: the option only displays the result in Stata output. In any other situations (*beta* or *ct* methods, no CIs, or assuming $K \perp p$), the option performs the test in addition to estimation and potential inference. In both cases, the number of bootstrap repetitions used for the test is the same as the one specified by option repbootstrap. When the user wants to test the binomial assumption, we recommend always doing so combined with inference using HR method in the general case (namely, without assuming independence between K and p): together with the test, it will give estimation and confidence intervals from np method virtually for free. The option is only available in unconditional analyses.¹⁵
- repct sets the number S of draws used to estimate θ_N^{ra} in CT's correction. Its argument needs to be a positive integer. The default value is 50.
- atkinson allows the user to specify the parameter b of the Atkinson index. Its argument has to be a real in (0,1). The default value is 0.5; it is the only one that ensures the symmetry property for the Atkinson index (i.e., the index does not change when swapping the minority/majority labels).

1.3.4 Saved results

The objects saved by **segregsmall** depend on the options, in particular, whether the analysis is unconditional or conditional. They can be gathered into three types of information about: (i) the data included in the analysis, (ii) the method and assumptions used, (iii) the estimation and inference results.

In this section, we list the objects saved in e() by the command and detail their contents when they relate to estimation and inference results. The remaining objects have self-explanatory names and are described in the help page of the segregsmall command.

Data included in the analysis Below, names with prefix I_ denote dummy variables equal to 1 if what follows is true, 0 otherwise. Objects stored in unconditional analyses are printed in black. *Additional* objects stored in conditional analyses are displayed in gray. The superscript *u indicates that the objects are only relevant and saved for unit-level covariates, the superscript *i for the individual-level covariates.

¹⁵The test of the assumption type by type can be done manually by restricting the sample used through the options [if][in].

Scalars:

```
e(I_withsingle)
                                                    e(nb_individuals)
                                                   e(nb_minority_individuals)
   e(I_excludingsinglepertype)
   e(I_unit_level_characteristic)
                                                   e(prop_minority_hat)
   e(nb_types)
                                                   e(K_max)
                                                   e(nb_K_with_obs)
   e(nb_units_total)
   e(nb units single)
                                                   e(nb cells studied sum across type)<sup>*i</sup>
                                                   e(nb_single_cells_sum_across_type)*i
   e(nb_units_studied)
Matrices:
   e(list_K_with_obs)
                                                   e(summary_info_data_per_type)
   e(type_frequencies)
                                                   e(nb_units_studied_per_type)<sup>*u</sup>
                                                    e(nb_cells_studied_per_type)*i
   e(type_probabilities)
Method used
Scalars:
```

```
e(I_method_np)e(I_noci)e(I_method_beta)e(nb_bootstrap_repetition)e(I_method_ct)e(specified_level)e(I_conditional)e(I_testbinomial)e(I_unit_level_characteristic)e(nb_ct_repetition)e(I_hyp_independenceKp)e(b_atkinson)
```

Estimation and inference Objects relative to unconditional analyses are in black (left-hand column); those relative to conditional analyses are in gray (right-hand column). Superscripts *np and *beta indicate that objects are only relevant and saved with np and beta method. Scalars:

The matrices whose name includes estimates_ci store the results of estimation and possible inference. The content of e(estimates_ci) varies with the method used but its structure remains similar. Each row corresponds to an index.

With *beta* method, ten rows represent the two possible aggregated indices θ_u (unit-level weights) and θ_i (individual-level weights), when K is considered as random, for each of the five indices (Duncan, Theil, Atkinson, Coworker, and Gini). For each possible index \times weights, the columns store the estimated index using R method with a Beta distribution restriction on F_p , and asymptotic confidence intervals at the traditional 90%, 95%, and 99% levels (plus the one specified by **level** if any).

With np method, the rows are identical, but there are only eight parameters since the Gini indices are absent. For each possible index × weights, e(estimates_ci)'s columns save: the estimated bounds $\hat{\underline{\theta}}_u$ and $\hat{\overline{\theta}}_u$ for unit-level weights (or $\hat{\underline{\theta}}_i$ and $\hat{\overline{\theta}}_i$ for individual-level weights); a dummy variable equal to 1 if the confidence interval used is the boundary-case interval and 0 for the interior-case; the resulting asymptotic CI at the classical 90%, 95%, and 99% levels (plus the one specified by level if any).¹⁶

 $^{^{16}{\}rm The}$ boundary/interior CIs were discussed briefly in Section 1.2.2 §Inference. We refer to the original paper HR for further details.

In conditional analyses, either with unit- or individual/position-level covariates, the matrices e(estimates_ci_aggregated) and e(estimates_ci_type_#) store exactly the same information as e(estimates_ci): the former for the aggregated conditional index $\theta_{0,u}^{\text{cond}}$ or $\theta_{0,i}^{\text{cond}}$, the latter for the index conditional on a given type #, that is θ_{0z} with unit-level characteristics or θ_{0w} with individual/position-level characteristics (# ranges from z = 1 to \overline{Z} or w = 1 to \overline{W}).

With ct method, five rows correspond respectively to the Duncan, the Theil, the Atkinson, the Coworker, and the Gini indices. In columns: the plug-in estimate $\hat{\theta}_N$ of the naive index θ_N ;¹⁷ the estimated index under random allocation $\hat{\theta}_N^{\text{ra}}$; the estimated CT-corrected index $\hat{\theta}_{CT}$; the empirical standard deviation of the draws $(\hat{\theta}_{N,s}^{\text{ra}})_{s=1,\dots,S}$ under random allocation; the "standardized score" originally proposed by Cortese et al. (1976), namely $(\hat{\theta}_N - \hat{\theta}_N^{\text{ra}})$ divided by that standard deviation; the empirical quantiles of $(\hat{\theta}_{N,s}^{\text{ra}})_s$ at the orders: 0.01, 0.05, 0.10, 0.90, 0.95, 0.99 (and also τ and $1 - \tau$, with τ the argument of option level provided this option is used).

 $e(I_constrained_case)$ is a dummy equal to 1 in the constrained case, 0 otherwise. As discussed in Section 1.2.2, with random unit size, it requires to be in the constrained case for each size $k \in \mathcal{K}$. In this case, np method yields point-estimates for all indices. $e(I_constrained_case)$ is identical in conditional analyses. The dummy is equal to 1 provided we are in the constrained case for each type. Otherwise, $\theta_{0,u}^{cond}$ and $\theta_{0,i}^{cond}$ are only partially-identified with np method.

 $e(test_binomial_results)$ is stored when the test of the binomial assumption is performed (see option testbinomial). It is a row vector whose first element saves the value of the test statistic LR_n and the second the p-value of the test where the null hypothesis is the binomial assumption.

np and beta methods save $e(info_distribution_of_p)$ in unconditional analyses.¹⁸ This matrix contains the information learned about the distribution of p in the estimation. In the general case, without assuming $K \perp p$, it means the information as regards the conditional distributions F_p^k , for each $k \in \mathcal{K}$. With option independencekp, it is about the unconditional distribution F_p .

With beta option, all the $(F_p^k)_{k\in\mathcal{K}}$ (or F_p when assuming $K \perp p$) are supposed to follow a Beta distribution. In the general case, $e(info_distribution_of_p)$ is a matrix with $|\mathcal{K}|$ rows. Each row is associated with a size k and the columns report: the size k; the number of units of size k in the data used, i.e. $\sum_{i=1}^{n} \mathbb{1}\{K_i = k\}$; the latter quantity expressed as a proportion over the n units studied; the number of components of the Beta mixture considered (that is 1); and the maximum likelihood estimators $\hat{\alpha}_1$ and $\hat{\beta}_1$ of the two shape parameters characterizing the Beta distribution assumed for F_p^k . In the case where $K \perp p$ is supposed, the matrix $e(info_distribution_of_p)$ is similar but consists of a single row as only one estimation is done pooling all units together. It contains the maximal size \overline{K} , the number of units n used for the estimation, and the estimates of the parameters that characterize the Beta distribution assumed for F_p .

With np option, the structure of $e(info_distribution_of_p)$ is more involved for the approach is nonparametric. Without the restriction $K \perp p$, it contains $3 \times \overline{K}$ rows and should be read by blocks of three rows. The k-th block concerns F_p^k . The first line shows some general information, namely the size k, the number of units of size k, and the

¹⁷Appendix 1.B presents the explicit expression of $\hat{\theta}_N$ for each index.

¹⁸In conditional analyses, the information can be retrieved manually for each type by restricting the sample used thanks to the options [if][in].

proportion of such units within the data used (as in *beta* method). The crucial element is displayed in the fourth column and consists of a dummy variable equal to 1 if we are in the constrained case for F_p^k , that is $\widehat{m} \in \partial \mathcal{M}$ conditional on K = k. In this case, despite the nonparametric approach, the constrained maximum likelihood estimation yields an estimate \widehat{F} of F_p^k which turns out to be a discrete distribution with at most $\lfloor (k+1)/2 \rfloor + 1$ support points (see Section 1.2.2 §Estimation). In this situation, the fifth column of the first row, within the three-row block, indicates the number of support points of F and the two following rows characterize F by reporting its support points and the corresponding probabilities. In the unconstrained case, the dummy is 0, and the last two rows, within the three-row block, are empty as there is no estimate of F_p^k then. When assuming $K \perp p$, the matrix e(info_distribution_of_p) is analogous but is made of a single three-row block as it only deals with the unconditional distribution F_p . In this case (see Section 1.2.2) Assuming independence between K and p, the estimation uses the first \overline{K} moments of F_p . It is likely to fall in the constrained case since \overline{K} will exceed 10 in most applications, a size above which simulations reveal that the probability of being in the constrained case is close to one even with large sample sizes n.

e(info_distribution_of_p) is interesting because virtually any segregation index is a functional of the distribution F_p (of the conditional distributions $(F_p^k)_{k\in\mathcal{K}}$ in general when taking into account the randomness of K). Consequently, an estimate of F_p (respectively of the $(F_p^k)_{k\in\mathcal{K}}$) enables to recover any other personalized segregation index.

1.3.5 Execution time

The computation times reported below are average over 50 repetitions on a desktop computer run under Windows 10 Enterprise with an Intel(R) Core(TM) i5-6600 CPU 3.30GHz processor (RAM 16 Go).¹⁹ The operations of segregsmall can be decomposed into a preparation stage and a stage devoted to estimation and inference.

Preparation stage The preparation stage is common to the three methods and reshapes the dataset. Its execution time is quick compared to the whole command and increases in the number n of units. For instance, with unit-level datasets, for K taking values in $\mathcal{K} = [5, 15]$, it lasts around 0.06 second with n = 1,000, and 0.99 second with n =300,000. In conditional analyses, the execution time is approximately multiplied by the number of types: for example, 6.03 seconds for 5 types and 9.17 seconds for 10 types, with $\mathcal{K} = [5, 15]$ and n = 300,000. With individual-level datasets, the preparation stage is longer since it is necessary first to form the units. With $\mathcal{K} = [5, 15]$, it takes 0.24 seconds with 1,000 units and 9.99 seconds with 300,000 units.

Estimation and inference stage The subsequent operations depend on the method used. The central brick of *np* and *beta* methods is the estimation of the indices for a given dataset (original or bootstrapped). The construction of CIs repeats the operation for each bootstrapped dataset. The execution time is thus more or less linear in the number of bootstrap repetitions (fixed by option repbootstrap). *ct* method requires reshuffling the data under the randomness benchmark, hence an execution time broadly linear in the number of draws (controlled by option repct). Table 1.1 illustrates this dependence for

¹⁹Execution times depend on various parameters. Rather than giving precise time guarantees, this section aims at outlining the main operations of segregsmall, their duration, and the dependence with respect to the sample size n, the unit size K, and the number \overline{Z} or \overline{W} of types in conditional analyses.

np and beta methods as well as the effect of option conditional. Regarding the latter, for all three methods, the same operations as in unconditional analyses are done for each type (see Section 1.2.6). As a consequence, the execution time of segregsmall is roughly linear in the number of types included in the analysis.

Table 1.1: Execution time in seconds. Setting: unit-level datasets, n = 300,000, $\mathcal{K} = [5, 15]$, 200 bootstrap replications, 5 types with covariates at unit-level for the conditional analysis.

Analysis	Confidence intervals	beta method	np method
unconditional	no	3.2	$1.3 \\ 176.9 \\ 906.8$
unconditional	yes	374.2	
conditional	yes	1870.8	

As highlighted by Table 1.2, the number n of units has a minor impact, mainly through the preparation stage.

Table 1.2: Execution time in seconds. Setting: unit-level datasets, $\mathcal{K} = [5, 15]$, options independencekp and noci for np and beta methods, 50 draws (default) for ct method.

Sample size n	beta method	np method	ct method
1,000	0.30	2.39	0.51
10,000	0.34	2.60	0.80
50,000	0.46	2.19	0.88
100,000	0.67	2.68	1.13

The primary determinant of the computation time is the unit sizes: both the number of distinct values of the support \mathcal{K} and the magnitude of K, as shown by Table 1.3. With ct method, the execution time quickly increases with the magnitude of K while the increase is moderate for np method and even lighter for *beta* method.²⁰

1.4 Example

We use the command to measure workplace segregation between natives and foreigners in France (see D'Haultfœuille and Rathelot (2017) for details about the context). A large share of workers is employed in small establishments. This section shows the importance of correcting the small-unit bias, which may lead to erroneous economic conclusions.

The data used is the 2007 *Déclarations Annuelles des Données Sociales* (DADS), French data linking workers to their employer. Data are exhaustive in the private sector (1.77 million establishments). In the application, we use the 1.04 million establishments that have between 2 and 25 employees. The minority group consists of individuals born outside of France and with the nationality of a country outside Europe. The overall proportion of minority individuals is 4.1% in the sample studied. Figure 1.1 shows the estimates of workplace segregation by firm size for the Duncan, the Theil, the Atkinson

²⁰Remark however that the execution times reported in Table 1.3 include the draws under random allocation for ct method whereas only estimation is performed for np and beta methods.

Support \mathcal{K} of K	beta method	np method	ct method
[1, 9]	0.28	0.99	0.23
[10, 19]	0.31	2.26	0.57
[20, 29]	0.26	5.10	2.16
[30, 39]	0.31	7.45	6.26
[40, 49]	0.36	8.20	15.1
[50, 59]	0.42	12.7	30.7
[60, 69]	0.51	11.0	56.6
[70, 79]	0.59	15.5	93.1
[80, 89]	0.70	22.7	150.3
[90, 99]	0.81	24.1	232.0
[100, 109]	0.93	26.3	332.1

Table 1.3: Execution time in seconds. Setting: unit-level datasets, for each \mathcal{K} , n = 10,000 (except 9,000 for the first row) – 1,000 units per distinct size, option **noci** for np and *beta* methods, 50 draws (default) for ct method.

(with parameter b = 0.5), and the Coworker indices. The Gini index does not satisfy the conditions required by the nonparametric method of HR and is thus not displayed.²¹



The distinct methods of the package are used: the estimated bounds $\underline{\hat{\theta}}$ and $\overline{\hat{\theta}}$ by np²¹Appendix 1.D presents some supplementary material for the example, including Figure 1.2 with the estimated Gini index by the *beta* method.

method on θ ("np bounds"); the 95%-level confidence interval for this parameter using the modified bootstrap procedure of np method, with the default 200 bootstrap iterations ("np CI"); the point-estimate $\hat{\theta}$ by *beta* method ("beta"); the plug-in estimate $\hat{\theta}_N$ of the naive index ("naive"); the estimated CT-corrected index $\hat{\theta}_{CT}$ using ct method with the default 50 draws under random allocation ("ct").

Figure 1.1 shows that the naive indices overestimate the actual level of segregation: they are almost always above the confidence interval obtained by np method (except for the Atkinson index with $K \in \{7, 8\}$). This bias decreases with the size of the units. For the Duncan, the Theil, and the Atkinson indices, the estimated identification interval for θ quickly becomes informative for $K \geq 5$ and reduces to a singleton for $K \geq 9$. When the unit size is larger than 1, the estimated bounds of np methods boil down to a point-estimate for the Coworker index (see discussion in Section 1.2.2).

The point-estimate $\hat{\theta}$ using *beta* method is within the identification bounds of HR for the Duncan, the Theil, and the Atkinson indices but is below HR's confidence intervals for the Coworker index. The CT-corrected measure $\hat{\theta}_{CT}$ underestimates the Duncan and Theil indices, being always below the *np* method's confidence interval. $\hat{\theta}_{CT}$ lies within the confidence interval and is quite close to the estimated identification set of θ for the Atkinson and Coworker indices.

Interestingly, the naive indices exhibit a stronger negative relationship between segregation levels and unit size than the corrected ones. Neglecting the small-unit bias would produce a statistical artifact as the magnitude of the bias decreases with K and therefore would support a negative correlation while it may not be so. On the contrary, the probability-based indices θ that account for the small-unit bias are able to address this question (see Section 5 of HR for further details).

Finally, we report below the Stata output obtained with the segregsmall command for np method and with option testbinomial. Appendix 1.D displays the output associated with beta and ct methods. Compared to the analyses of Figure 1.1 (K by K), the estimation is performed over the entire sample of units ($\mathcal{K} = [2, 25]$) in this output without assuming $K \perp p$. As detailed in Section 1.3.3, the test of the binomial assumption is automatically performed and saved in this configuration; the option only displays the result in the Stata output. In this application, we cannot reject the binomial assumption at any standard level.

```
. segregsmall K X, format(unit) method(np) testbinomial
*** Construction of relevant databases for the analysis ***
*** Estimation and inference ***
Estimation - current unit size analyzed (out of 24 distinct sizes):
Preparation of bootstrap -
Bootstrap - current bootstrap iteration (out of 200):
Bounds for segregation indices using nonparametric (np) method:
Unconditional analysis
Number of units studied in the analysis: 1036840
(0 unit with a single individual are excluded from the analysis)
Number of individuals studied: 6178564
Proportion of minority (or reference) group: 4.1e-02
Assumption on dependence between K and p for estimation and inference: none
Inference: by bootstrap, 200 repetitions
Unconditional segregation indices:
  Index
           Weight-level
                         Lower bound
                                    Upper bound
                                                [95% Conf. Interval]
```

Duncan	unit	.58677	.82346	.57864	.8292	
Duncan	individual	.63061	.74808	.61966	.75742	
Theil	unit	.39246	.52092	.38901	.5251	
Theil	individual	.37937	.44251	.37558	.44732	
Atkinson	unit	.53907	.83164	.52638	.84667	
Atkinson	individual	.56948	.73299	.54977	.7537	
Coworker	unit	.37032	.37032	.3674	.37325	
Coworker	individual	.31356	.31356	.31084	.31629	
Test of binom (distribution	ial assumption (under the null o	HO: conditional N obtained by boots	oinomial distri strap, 200 repe	bution): titions)		
Result		value of test statistic p-		p-value	lue	
		1.53	98598	.23		

1.5 Conclusion

This paper presented the Stata segregsmall command that implements three methods (D'Haultfœuille and Rathelot (2017), Rathelot (2012), and Carrington and Troske (1997)) to measure segregation indices in settings when units (neighborhoods, firms, classes, etc.) contain few individuals. In such situations, naive indices overestimate the actual level of segregation and produce measures that are not comparable across settings or over time since the small-unit bias might vary. segregsmall enables social scientists to compute segregation indices in those cases and makes the HR nonparametric approach easy to use. It provides asymptotic confidence intervals for HR and R parameters. For all three methods, conditional indices can be estimated: they account for other covariates (either at unit- or individual/position-level) that may influence the allocation process of individuals into units and therefore measure "net" or "residual" segregation. HR and R methods can be used whatever the unit size to measure segregation as a departure from the relevant benchmark of randomness. Even with large units with above one hundred individuals, the parametric approach of R method remains quite affordable as regards computational requirements, even including inference by bootstrap.

References

- Allen, R., Burgess, S., Davidson, R. and Windmeijer, F. (2015), 'More reliable inference for the dissimilarity index of segregation', *The Econometrics Journal* 18(1), 40–66.
- Åslund, O. and Skans, O. N. (2009), 'How to measure segregation conditional on the distribution of covariates', *Journal of Population Economics* **22**(4), 971–981.
- Boisso, D., Hayes, K., Hirschberg, J. and Silber, J. (1994), 'Occupational segregation in the multidimensional case: decomposition and tests of significance', *Journal of Econometrics* 61(1), 161–171.
- Carrington, W. J. and Troske, K. R. (1997), 'On measuring segregation in samples with small units', Journal of Business & Economic Statistics 15(4), 402–409.
- Cortese, C. F., Falk, R. F. and Cohen, J. K. (1976), 'Further considerations on the methodological analysis of segregation indices', *American Sociological Review* pp. 630– 637.
- Dette, H. and Studden, W. J. (1997), The theory of canonical moments with applications in statistics, probability, and analysis, Vol. 338, John Wiley & Sons.
- D'Haultfœuille, X. and Rathelot, R. (2017), 'Measuring segregation on small units: A partial identification analysis', *Quantitative Economics* 8(1), 39–73. URL: http://dx.doi.org/10.3982/QE501
- Imbens, G. W. and Manski, C. F. (2004), 'Confidence intervals for partially identified parameters', *Econometrica* 72(6), 1845–1857.
- James, D. R. and Taeuber, K. E. (1985), 'Measures of segregation', Sociological Methodology 15, 1–32.
- Krein, M. and Nudel'man, A. (1977), 'The markov moment problem and extremal problems transl', *Mathematical Monographs* **50**.
- Massey, D. S. and Denton, N. A. (1988), 'The dimensions of residential segregation', Social forces 67(2), 281–315.
- Rathelot, R. (2012), 'Measuring segregation when units are small: A parametric approach', Journal of Business & Economic Statistics 30(4), 546–553.
 URL: http://www.tandfonline.com/doi/abs/10.1080/07350015.2012.707586
- Reardon, S. F. and Firebaugh, G. (2002), 'Measures of multigroup segregation', Sociological Methodology 32(1), 33–67.
- Stoye, J. (2009), 'More on confidence intervals for partially identified parameters', *Econometrica* 77(4), 1299–1315.
- Winship, C. (1977), 'A revaluation of indexes of residential segregation', Social Forces 55(4), 1058–1066.

Appendix 1.A Magnitude of the small-unit bias

Tables 1.4 and 1.5 present the replicated results of the simulation studies made by Cortese et al. (1976) and Carrington and Troske (1997) to illustrate the magnitude of the small-unit bias. For the Duncan index (Table 1.4) and the Gini index (Table 1.5), they report $\hat{\theta}_N^{\text{ra}}$ (see Section 1.2.4), namely the estimated expected value of the proportion-based or naive index under random allocation of the individuals across units (using in those simulations S = 500 replications). The simulations are performed for a fixed number n = 100 of units and several unit sizes K and minority proportions to assess the magnitude of the bias in these different settings.

Table 1.4: Estimated $\hat{\theta}_N^{\text{ra}}$ for the Duncan: expected value of the naive index under random allocation (that is, null polarization).

Minority share	K = 2	K = 5	K = 10	K = 20	K = 50	K = 100	K = 1,000
1%	0.95	0.93	0.90	0.82	0.61	0.37	0.13
2%	0.96	0.92	0.83	0.69	0.39	0.27	0.09
5%	0.94	0.80	0.63	0.39	0.26	0.18	0.06
10%	0.90	0.65	0.39	0.29	0.19	0.13	0.04
20%	0.81	0.42	0.31	0.22	0.14	0.10	0.03
30%	0.73	0.41	0.27	0.19	0.12	0.09	0.03
40%	0.68	0.35	0.26	0.18	0.11	0.08	0.03
50%	0.66	0.37	0.25	0.18	0.10	0.08	0.02

Table 1.5: Estimated $\hat{\theta}_N^{\text{ra}}$ for the Gini: expected value of the naive index under random allocation (that is, null polarization).

Minority share	K = 2	K = 5	K = 10	K = 20	K = 50	K = 100	K = 1,000
1%	0.99	0.96	0.92	0.83	0.69	0.52	0.18
2%	0.98	0.92	0.84	0.72	0.53	0.39	0.13
5%	0.95	0.83	0.69	0.54	0.33	0.25	0.08
10%	0.90	0.72	0.55	0.41	0.26	0.18	0.06
20%	0.80	0.59	0.43	0.31	0.20	0.14	0.04
30%	0.73	0.53	0.38	0.27	0.17	0.12	0.04
40%	0.68	0.49	0.36	0.25	0.16	0.11	0.04
50%	0.65	0.49	0.34	0.25	0.16	0.11	0.04

In the simulations, the actual level of segregation (understood in the randomness benchmark) is null since individuals are distributed across units at random, independently of belonging to the minority group. The large positive values obtained for $\hat{\theta}_N^{\rm ra}$ relative to the [0, 1] range of segregation indices indicate that naive indices overestimate the level of segregation. In various applications (a large share of firms have less than ten employees, school classes typically have between twenty and forty pupils, only surveys might be available for residential segregation yielding a small number of observed inhabitants by neighborhoods), the bias is likely to be non-negligible since the estimated $\hat{\theta}_N^{\rm ra}$ remain far from 0 even with moderate to large unit sizes, for instance, for the Duncan, 0.29 with K = 20 (respectively 0.19 with K = 50) for a minority group whose proportion represents 10% in the overall population. The simulations also evidence that the small-unit bias decreases with the unit sizes and with the share of the minority group in the population (more precisely, all else being equal, the closer the share of the minority group, or reference group, to one half, the smaller the bias).

Appendix 1.B Plug-in estimators of naive indices

We write below the expressions of the plug-in estimator $\hat{\theta}_N$ of the naive or proportion-based index $\theta_N = g\left(F_{X/K}, m_{01}\right)$ for the five classical indices we consider, namely the Duncan (a.k.a. dissimilarity index), the Theil (a.k.a. entropy or information theory segregation index), the Atkinson parameterized by $b \in (0, 1)$, the Coworker (a.k.a. normalized isolation or variance ratio index), and the Gini.

Let $w_i = K_i/(\sum_{j=1}^n K_j)$ the weight (relative size in terms of number of individuals) of unit *i* and $\pi = (\sum_{i=1}^n X_i)/(\sum_{i=1}^n K_i)$ the sample proportion of the minority group. The plug-in estimators $\hat{\theta}_N$ are

$$\begin{split} \widehat{\theta}_{N}^{D} &= \frac{1}{2\pi(1-\pi)} \sum_{i=1}^{n} w_{i} \left| \frac{X_{i}}{K_{i}} - \pi \right| \quad \text{(Duncan)}, \\ \widehat{\theta}_{N}^{T} &= 1 - \sum_{i=1}^{n} w_{i} \frac{(X_{i}/K_{i}) \ln\left\{ (X_{i}/K_{i})^{-1} \right\} + (1 - (X_{i}/K_{i})) \ln\left\{ (1 - (X_{i}/K_{i}))^{-1} \right\}}{\pi \ln\left\{\pi^{-1}\right\} + (1 - \pi) \ln\left\{ (1 - \pi)^{-1} \right\}} \quad \text{(Theil)}, \\ \widehat{\theta}_{N}^{A(b)} &= 1 - \pi^{\frac{-b}{1-b}} \left(1 - \pi \right)^{-1} \left[\sum_{i=1}^{n} w_{i} \left(1 - \frac{X_{i}}{K_{i}} \right)^{1-b} \left(\frac{X_{i}}{K_{i}} \right)^{b} \right]^{\frac{1}{1-b}} \quad \text{(Atkinson with } b \in (0, 1)), \\ \widehat{\theta}_{N}^{CW} &= \frac{1}{\pi(1-\pi)} \sum_{i=1}^{n} w_{i} \left(\frac{X_{i}}{K_{i}} - \pi \right)^{2} \quad \text{(Coworker)}, \\ \widehat{\theta}_{N}^{G} &= \frac{1}{2\pi(1-\pi)} \sum_{i=1}^{n} \sum_{j=1}^{n} w_{i} w_{j} \left| \frac{X_{i}}{K_{i}} - \frac{X_{j}}{K_{j}} \right| \quad \text{(Gini)}. \end{split}$$

We note that there have been debates as regards which index should be used. This constitutes another issue as regards the problem of quantifying segregation. Nonetheless, the opposition between evenness and randomness benchmarks is orthogonal and relevant for any segregation index. We focus on the latter in the article, and the command **segregsmall** returns those five classical indices, leaving the choice of the index and its axiomatic properties to users.

Remark that each $\hat{\theta}_N$ is exactly the corresponding segregation index as defined in James and Taeuber (1985) and Massey and Denton (1988). Indeed, in these works, there was no distinction between the estimand and the estimator: segregation indices were directly defined from the observed proportions $(X_i/K_i)_{i=1,...,n}$ (typically, the setting was a complete census of the tracts in a given city) instead of the distribution function $F_{X/K}$.

Appendix 1.C Indices in the parametric approach

We use here the same notation as in Section 1.2.3: $B(\cdot, \cdot)$ denotes the beta function, $B(x,y) = \int_0^1 u^{x-1}(1-u)^{y-1} du$; *c* the number of components of the beta mixture, and $v = (\alpha_j, \beta_j, \lambda_j)_{j=1,...,c}$ the vector of parameters with $(\alpha_j, \beta_j) \in \mathbb{R}^*_+ \times \mathbb{R}^*_+$ the two shape parameters of the *j*-th Beta distribution and $\lambda_j \in [0, 1]$ its weight $(\sum_{j=1}^c \lambda_j = 1)$. If B is a random variable distributed according to the mixture of Beta distributions characterized by v, we have

$$\mu(v) = \mathcal{E}(B) = \sum_{j=1}^{c} \lambda_j \frac{\alpha_j}{\alpha_j + \beta_j}.$$

Duncan index Let I(t; a, b) = B(t; a, b)/B(a, b) with $B(t; a, b) = \int_0^t u^{a-1}(1-u)^{b-1} du$ the incomplete beta function. Using B(a, b+1) = B(a, b)b/(a+b) and I(1-t; a, b) = 1-I(t; b, a), we obtain

$$\theta^{\rm D} = \frac{\mu(v) \sum_{j=1}^{c} \lambda_j I(\mu(v); \alpha_j, \beta_j) - \sum_{j=1}^{c} \lambda_j \alpha_j (\alpha_j + \beta_j)^{-1} I(\mu(v); \alpha_j + 1, \beta_j)}{\mu(v)(1 - \mu(v))}$$

Theil index To derive the expression of the Theil index as a function of v, we use that if $B \sim \text{Beta}(\alpha, \beta)$, then $1 - B \sim \text{Beta}(\beta, \alpha)$ and $\text{E}[B \ln(B)] = \alpha(\alpha + \beta)^{-1} \{\psi(\alpha + 1) - \psi(\alpha + \beta + 1)\}$, with ψ the digamma function. This yields

$$\theta^{\mathrm{T}} = 1 - \frac{\sum_{j=1}^{c} \lambda_j \left\{ \frac{\alpha_j}{\alpha_j + \beta_j} \psi(\alpha_j + 1) + \frac{\beta_j}{\alpha_j + \beta_j} \psi(\beta_j + 1) - \psi(\alpha_j + \beta_j + 1) \right\}}{\mu(v) \ln \{\mu(v)\} + (1 - \mu(v)) \ln \{1 - \mu(v)\}}$$

Atkinson index Let $\Gamma(t) = \int_0^{+\infty} u^{t-1} \exp(-u) du$ denote the gamma function. Using that $B(a,b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$ and $\Gamma(t+1) = t\Gamma(t)$, the Atkinson index satisfies, for any $b \in (0,1)$,

$$\theta^{\mathcal{A}(\mathbf{b})} = 1 - \frac{\mu(v)^{\frac{-b}{1-b}}}{1-\mu(v)} \left\{ \sum_{j=1}^{c} \lambda_j \frac{\Gamma(\alpha_j+b)\Gamma(\beta_j+1-b)}{\Gamma(\alpha_j)\Gamma(\beta_j)(\alpha_j+\beta_j)} \right\}^{\frac{1}{1-b}}.$$

Coworker index If $B \sim \text{Beta}(\alpha, \beta)$, then $E[B^2] = \alpha(\alpha+1)/\{(\alpha+\beta+1)(\alpha+\beta)\}$. This implies

$$\theta^{\mathrm{CW}} = \left\{ \sum_{j=1}^{c} \lambda_j \frac{\alpha_j(\alpha_j+1)}{(\alpha_j+\beta_j+1)(\alpha_j+\beta_j)} - \mu(v)^2 \right\} \left/ \left\{ \mu(v) - \mu(v)^2 \right\}.$$

Gini index Contrary to the previous indices, there is no closed-form expression for the Gini index under a mixture of Beta distributions for F_p because of the term $\int \{F_p(u)\}^2 du$. This quantity has to be approximated by numerical methods. The Gini index can only be written as

$$\theta^{\rm G} = \left[1 - \mu(v) - \int_0^1 \left\{ \int_0^u \sum_{j=1}^c \lambda_j \frac{t^{\alpha_j - 1} (1 - t)^{\beta_j - 1}}{B(\alpha_j, \beta_j)} \, \mathrm{d}t \right\}^2 \mathrm{d}u \right] \Big/ \left\{ \mu(v) - \mu(v)^2 \right\}.$$

Appendix 1.D Supplements to the example

Figure 1.2 is equivalent to the analysis displayed in Figure 1.1 for the Gini index. Because the Gini index does not satisfy Assumption 1.1, only the output of *beta* and *ct* methods are reported: the point-estimate $\hat{\theta}$ ("beta") and the 95%-level asymptotic confidence intervals

obtained by bootstrap ("beta CI") using *beta* method (with the default 200 bootstrap iterations); the plug-in estimate $\hat{\theta}_N$ of the naive or proportion-based index ("naive"); the estimated CT-corrected index $\hat{\theta}_{CT}$ using *ct* method with the default 50 draws under random allocation ("ct").

As with the Duncan, the Theil, the Atkinson, and the Coworker indices, Figure 1.2 illustrates the points discussed in Section 1.2 in the particular case of the Gini index. Regarding CT correction, there is no reason why θ_{CT} should be close to θ . For the Gini, the CT-corrected index happens to fall below the confidence interval on the probability-based index θ obtained by *beta* method.



Figure 1.2: Gini index by firm size.

We report below the Stata output obtained with *beta* and *ct* methods. These estimations are done over the whole sample of units ($\mathcal{K} = [2, 25]$). As an illustration, the option independencekp is used for *beta* method.

```
. segregsmall K X, format(unit) method(beta) independencekp repb(400) level(0.98)
*** Construction of relevant databases for the analysis ***
*** Estimation and inference ***
Estimation - K and p assumed independent: units are merged (maximal size = 25)
Bootstrap - current bootstrap iteration (out of 400):
Estimates for segregation indices using parametric (beta) method:
Unconditional analysis
Number of units studied in the analysis: 1036840
(0 unit with a single individual are excluded from the analysis)
Number of individuals studied: 6178564
Proportion of minority (or reference) group: 4.1e-02
```

Index	Weight-level	Point-estimate	[98% Conf.	Interval]
Duncan	unit	.75967	.75777	.76129
Duncan	individual	.75967	.75777	.76129
Theil	unit	. 43393	.43098	.43639
Theil	individual	. 43393	.43098	.43639
Atkinson	unit	.76516	.76254	.76741
Atkinson	individual	.76516	.76254	.76741
Coworker	unit	. 2795	.27604	.28258
Coworker	individual	.2795	.27604	.28258
Gini	unit	.89272	.89135	.89388
Gini	individual	.89272	.89135	.89388
		1		

Assumption on dependence between K and p for estimation and inference: independence Inference: by bootstrap, 400 repetitions Unconditional segregation indices:

. segregsmall K X, format(unit) method(ct) repct(100)
*** Construction of relevant databases for the analysis ***
*** Estimation and correction ***
CT-correction - current random allocation iteration x10 (out of 100):
.....+
Estimates for segregation indices using CT-correction (ct) method:

```
Unconditional analysis
```

```
Number of units studied in the analysis: 1036840
```

(0 unit with a single individual are excluded from the analysis)

```
Number of individuals studied: 6178564
```

Proportion of minority (or reference) group: 4.1e-02 No inference for naive and CT-corrected indices

CT-correction is made using 100 draws under random allocation (u.r.a.)

Unconditional segregation indices:

Index	Weight-level	Naive	Expected u.r.a.	CT-corrected
Duncan	n.a.	.85864	.71364	.50634
Theil	n.a.	.57585	.35113	.34632
Atkinson	n.a.	.90735	.75392	.62349
Coworker	n.a.	.41953	.16779	.3025
Gini	n.a.	.94481	.832	.67147

Chapter 2

Measures of several dimensions of residential segregation in France between 1968 and 2019 based on the Labor Force Survey clusters

The first chapter presented the Stata command segregsmall. This second chapter applies it to measure residential segregation in France. The analysis is done over the period 1968-2019 to study temporal evolution. I also compare the magnitude of segregation across various dimensions, namely different ways (using demographic and socioeconomic individual variables) to specify the minority and majority groups.

I thank INSEE, "le réseau Quetelet-Progedo Diffusion", and "le Comité du Secret Statistique" for the access to the datasets used in this project: Enquête Emploi 1968-2002, Enquête Emploi en continu (version FPR) 2003-2019.

Abstract This article takes advantage of the sampling scheme of the French Labor Force Survey, which draws clusters of around thirty adjacent housing, to study several dimensions of residential segregation in France over 1968-2019. Such clusters form relevant neighborhoods to study residential segregation provided the indices account for the small-unit bias so that they can be compared over time or across different dimensions of segregation (French versus non-French people, jobseekers versus employed, college graduates versus non-graduate, white-collar versus blue-collar workers, etc.). Applying the methodology developed in D'Haultfœuille and Rathelot (Quantitative Economics, 2017) and Rathelot (Journal of Business and Economic Statistics, 2012), we estimate annual segregation indices for different specifications of the "minority" and "majority" groups, aiming to quantify several dimensions of residential segregation and compare them. The results suggest two main conclusions. First, whatever the dimension under study (ethnicity, immigrant, nationality, occupational category, labor market status, education), the estimated indices do not reveal significant evolution over time: within each dimension, the magnitude of residential segregation has remained globally constant for the past decades. Second, they reveal the magnitude of segregation differs across the different dimensions according to the following decreasing ranking: nationality, ethnicity (using as a proxy parents' country of birth) and being an immigrant, social status (occupational category or college education), and labor market status (unemployed or employed). A conditional analysis, separating neighborhoods that belong to urban areas of 200,000 inhabitants or more from neighborhoods belonging to smaller urban areas, complements the unconditional analysis.
2.1 Introduction

This article provides quantitative measures of residential segregation in France over the past decades, between 1968 and 2019, for several socioeconomic or demographic dimensions, namely several ways of splitting the population living in France into two groups, a so-called minority or reference group and the other majority group. For a given split, at a given date, a residential segregation index for metropolitan France quantifies the extent to which minority individuals are concentrated in some neighborhoods instead of being uniformly spread across the residential areas of metropolitan France. We use the clusters of about 30 adjacent housing from the Labor Force Survey as the studied neighborhoods. For each year, we estimate residential segregation indices at the scale of the whole metropolitan French territory. Such estimates enable us to study the evolution of residential segregation over time and compare the magnitude of segregation for different definitions of minority groups. They could also be used as basic bricks to study the causes of residential segregation.

The outline of the article is as follows. The rest of this introductory section briefly presents the data used, the notion of residential segregation, and the methods we apply to measure it while addressing the so-called small-unit bias. It then defines what we mean by a dimension of residential segregation, discusses the related literature, and states the main findings. Section 2.2 describes the data we use in more detail and how we define the various dimensions, that is, the specifications of a minority group and of a majority group according to socioeconomic or demographic individual covariates. Section 2.3 formally presents segregation indices as statistical parameters of interest and the estimation and inference methods we use. It also shows several robustness checks. Finally, Section 2.4 presents the estimated segregation indices over time and for the different dimensions.

Data and neighborhoods/residential areas

From a population made of two specified groups (minority and majority) spread across units (neighborhoods, workplaces, schools, etc.), a segregation index seeks to measure the extent to which individuals from the minority group are concentrated in some units more than in others. Consequently, the notion depends on the units considered, for instance, school segregation with schools or residential segregation with neighborhoods.

Defining neighborhoods can be challenging. Here, we take advantage of the areal sampling scheme of the French Labor Force Survey (LFS). From an exhaustive sample frame of the ordinary housing¹ in France based on housing tax, the sampling randomly draws clusters ("grappes") of around 20 to 40 adjacent housing and surveys the people who live there. From 1968 to 2002, the survey was done annually, at a particular year period (typically during the first semester in March). Since 2003, the French LFS has become the "Enquête Emploi en continu," and the survey has been performed throughout the year with selected clusters surveyed every quarter during six quarters, one and a half years.² After being surveyed for six quarters, a cluster is dropped, and another is drawn

¹The population of study of the LFS is almost all *residents* in France, not the sub-population of French individuals; almost because the survey concerns inhabitants who, at the time of the survey, are at least 15 years old and live in *ordinary* housing, which excludes individuals living in communities such as retirement homes, hospitals, barracks, monasteries, etc. Children younger than 15 years old are not concerned by the survey, although it is possible to identify them and have some basic information.

²Note that the survey does not follow inhabitants who move; it encompasses the *current* inhabitants of the cluster.

to replace it. In fact, the panel is rotative in the sense that, every quarter, one-sixth of the clusters entering the survey are renewed.

Despite the variation in their physical delimitation, according to population densities, these clusters of around 30 adjacent housing form relevant *units* or *neighborhoods* (henceforth, the two terms are used as synonyms) to study residential segregation. They do capture individuals whose places of residence are close together.

Residential segregation

As already explained, for a given partition of the population into two groups, a minority and a majority group, the notion of *residential segregation* relates to the concentration of minority individuals in some neighborhoods more than in others. Remark that, following the literature, we use the term segregation as a neutral term describing that concentration independently of the causes behind such a phenomenon. Measuring the level of segregation is a necessary first step to investigating its causes.

Qualitatively, there is some segregation whenever the minority individuals are not randomly distributed across the units; in other words, when the allocation process systematically leads to the concentration of minority individuals in some specific units: minority individuals are over-represented in some neighborhoods relative to their overall proportion in the population and under-represented in some others.

This suggests a natural way to quantify segregation by considering the empirical proportions $(X_i/K_i)_{i=1,...,n}$ where, for each unit $i \in \{1, ..., n\}$ among the *n* units surveyed, X_i (respectively K_i) is the number of minority individuals (respectively the total number of individuals) living in unit *i*. Intuitively, the more variation across those proportions, the higher residential segregation. Indeed, suppose that a large part of those proportions is equal to either 0 or 1, it means that most neighborhoods have only majority individuals or only minority individuals. On the contrary, if there is little variation in the proportions X_i/K_i across the units $i \in \{1, ..., n\}$, it means that minority individuals are more or less uniformly distributed across the neighborhoods. Consequently, a natural measure of segregation is an inequality index on the proportions $(X_i/K_i)_{i=1,...,n}$.

Small-unit bias

Although intuitive, considering the variations in the observed proportions is problematic when the units are small, namely when the number K of individuals per unit is small. In that case, the natural measure based on observed proportions $(X_i/K_i)_{i=1,...,n}$ over-estimates the level of segregation; this problem faced by naive indices is known as the small-unit bias. The issue is that with small K, the variation in the empirical shares $(X_i/K_i)_{i=1,...,n}$ could simply be due to small-sample variability, instead of reflecting a systematic concentration of minority individuals in some units more than in others; the smaller K, the more severe the bias.

Let forget temporary residential segregation and imagine a simple example to convey the idea. We consider a population of employees, half women, half men, and we study workplace segregation between gender: do women and men tend to work in the same workplaces (which correspond here to the units) or, on the contrary, do they concentrate into distinct workplaces? Assume that there is no segregation in the sense that, for each workplace, the probability that an arbitrary employee working there is a woman is equal to 1/2. In other words, women and men are allocated randomly across workplaces. Nevertheless, as long as the sizes of workplaces are small, the observed proportions of women and men will not be precisely equal to the theoretical 50-50% repartition owing to statistical fluctuations. In an extreme case with two employees per workplace, although women and men are allocated randomly, on average, one-fourth of the workplaces will have only women and another one-fourth only men.

Thus, in such situations, naive measures of segregation are upward-biased. This fact is problematic in itself to quantify segregation. Furthermore, it impedes reliable comparisons over time or across countries since the bias may evolve.

Significance of the bias in our application

We could wonder how significant the bias in our setting is since the units contain several dozens of individuals. Simulation studies from Cortese et al. (1976) and Carrington and Troske (1997) show that, even with K of the order of 50 individuals, the bias is substantial. In addition to K, all else equal, the magnitude of the bias decreases with the overall share of the minority group in the population: the closer to 1/2, the lower the bias. However, in several analyses performed below, when considering, for instance, immigrants or foreigners to define the minority group, that share is well below 1/2. To quote a single number from those simulation studies, with a minority share of 10% of population and K = 50, the estimated bias of the classical Duncan or dissimilarity index is 0.19, which is almost one-fifth of the range of the index, normalized between 0 (no segregation) to 1 (complete segregation).

Such magnitude of the bias, if not higher, can be expected in our application. Indeed, pooling all years 1968-2019 together,³ the number of adults (individuals aged 18 or more) by unit has the following descriptive statistics: a mean of 45 with a standard deviation of 33, a median of 36 and respective quantiles at 10%, 25%, 75% and 90% equal to 15, 23, 61 and 92.⁴

Above all, the small-unit bias impedes reliable comparisons of segregation measures over time, across countries, or different dimensions/specifications of minority and majority groups. Indeed, the size of units might change, thus modifying the magnitude of the bias. It is indeed the case for our data as the LFS's sampling scheme has evolved toward thinner clusters over time: the average number (median in parenthesis) of adults by cluster is 73 (68) between 1968 and 1979, 56 (58) for the decade 1980-1989, 42 (37) between 1990 and 2002, and 28 (28) for 2003-2019. Consequently, naive segregation measures are likely to exhibit an artificial increase of segregation over time since the small-unit bias is more acute as K gets smaller.

All in all, when measuring residential segregation with units based on the clusters of the LFS, it is vital to correct for the small-unit bias.

Methodology

Two main approaches have been proposed in the literature to do so. A first one (Cortese et al. (1976), Carrington and Troske (1997), Allen et al. (2015)) starts from the proportionbased indices (variations of the $(X_i/K_i)_{i=1,...,n}$) and corrects them by subtracting an

 $^{^{3}}$ As explained above, since 2003, the survey has been done four times a year, every quarter. To maintain annual regularity and study evolution over time, we restrict the analysis to one quarter; we choose the first one as, before 2003, the survey was generally conducted in March.

⁴Besides, in the following analyses, the number of individuals K actually included in the population of interest can even be smaller since we do not always observe the relevant covariates defining group membership for all individuals.

estimate of the bias.

Another approach (Rathelot (2012), D'Haultfœuille and Rathelot (2017)) directly defines probability-based indices immune to the small-unit bias. The units are modeled as *i.i.d.*, and the idea is to introduce, for each unit *i*, the unobserved variable p_i . p_i is the conditional probability that an arbitrary individual of unit *i*, knowing that he or she lives in that neighborhood *i*, belongs to the minority group.⁵ Remark that X_i/K_i is a natural estimator of p_i , but is consistent only when K_i goes to infinity. As already explained, the smaller K_i , the noisier the estimator, which implies variations in the proportions $(X_i/K_i)_{i=1,...,n}$ and small-unit bias.

In contrast, probability-based indices are defined from the variations of the probabilities p_i . If they are all equal, it exactly says that the minority individuals are distributed randomly across the neighborhoods: the probability that an arbitrary individual belongs to the minority group conditional on living in neighborhood *i* is equal to the unconditional probability (the overall share of the minority group in the population). Thinking in terms of information, this means that knowing the place of residence of someone does not give any clue to predict its membership to the minority. This is how null segregation is defined. On the contrary, a situation where all the probabilities p_i are equal to either 0 or 1 characterizes complete segregation. Formally, as detailed in Section 2.3, the indices are defined as functionals of the distribution of p.

The opposition between proportion-based and probability-based indices can be exposed as an opposition between *evenness* and *randomness* benchmarks. Following Winship (1977), it is noteworthy to stress that both benchmarks can be of interest, although for the question of the measure of segregation, we think the randomness one should be favored.

The evenness benchmark considers the realized proportions X_i/K_i . Suppose one is interested in the consequences or effects of segregation. In that case, those proportions are probably the relevant objects of interest since the impact will depend on the actual allocation of the population. In contrast, when the issue is to quantify the level of segregation, namely to which extent minority individuals tend to be concentrated in some neighborhoods more than others, the interest lies in the underlying allocation process, and the relevant quantities are arguably the probabilities p_i . We consider the randomness benchmark in the rest of this article.

For the estimation of and inference on the probability-based indices, we use the methodology developed in Rathelot (2012) and D'Haultfœuille and Rathelot (2017) and implemented in the Stata package segregsmall (D'Haultfœuille et al., 2021).

Orthogonal to the distinction between probability- and proportion-based indices, another debate regarding segregation measure concerns the choice of the mapping from either the proportions $(X_i/K_i)_i$ or the probabilities $(p_i)_i$ to a normalized segregation index in [0, 1]. We present below measures for five classical indices: the Duncan (also known as dissimilarity index), the Theil (a.k.a. entropy or information theory segregation index), the Atkinson (with parameter b = 0.5), the Coworker (a.k.a. normalized isolation or variance-ratio index), and the Gini.

Dimensions of segregation

By different *dimensions* of segregation, we understand different ways to define the minority and the majority groups under study. We consider several socioeconomic and demographic

⁵Such a probability is properly defined under the assumption: $X | K, p \sim \text{Binomial}(K, p)$, where we drop the subscript *i* to lighten notation as the variables are modeled as *i.i.d.*.

variables yielding six main dimensions of the analysis, depending on the type of individual covariates used to classify minority and majority individuals: (i) labor market positions, (ii) social/occupational categories, (iii) levels of education, (iv) ethnicity proxied by the country of birth of an individual's mother, (v) origins in the sense of immigrants or non-immigrants; (vi) nationalities.

A given dimension, that is, a type of individual covariates used to distinguish minority and majority, is susceptible to include different specifications of minority and majority groups. For instance, for the dimension of nationality, one analysis can be done with European nationalities defining the majority group and non-European nationalities as the minority one; another with French individuals as the majority group and individuals with Algerian, Moroccan, or Tunisian nationalities as the minority group.

That example raises the question of how to proceed when the specifications entail a partition of the overall population in more than two groups. It is the case in the latter specification where we can consider two minority and one majority groups:⁶ group M (standing for the majority) made of French, group m_1 (standing for minority) made of Algerians, Moroccans, or Tunisians, and group m_2 made of other foreign nationalities. Another example is the labor market dimension, where we can distinguish the group m_1 of jobseekers, the group M of employed people, and the group m_2 of inactive people.

Remark that the labels of *minority* or *majority* groups are conventional. First, all the indices we consider are invariant to the label-switching of the groups.⁷ Second, rather than the actual proportions in the population, these choices are guided by the sub-populations we want to study and contrast: we are interested in quantifying the concentration of the m_1 -type individuals in some units only compared to the distribution of the M-type individuals; the other group m_2 being of less or no interest. This explains the conventional choice in the literature since Massey and Denton (1988):⁸ drop group m_2 and define the studied population as the minority group m_1 and the majority group M to estimate segregation indices in a two-group setting. Compared to regrouping m_2 and M together, this choice is sensible by focusing on the concentration of the minority group of interest m_1 compared to a benchmark/reference group M. We follow this specification in all our analyses. Remark that dropping individuals from group m_2 leads to smaller unit sizes K; hence, addressing the small-unit bias is crucial.

Finally, we underscore that our use of the term *dimension* differs from the seminal article "The Dimensions of Residential Segregation" by Massey and Denton (1988) (MD) where the authors define residential segregation as a multidimensional phenomenon with five dimensions:

- 1. Evenness "refers to the differential distribution of two social groups among areal units in a city. A minority group is said to be segregated if it is unevenly distributed over areal units (Blau 1977)" (MD, page 283, quoting Blau (1977));
- Exposure "refers to the degree of potential contact, or the possibility of interaction, between minority and majority group members within geographic areas of a city" (MD, page 287);

⁶The following discussion would remain unchanged if we have more than two minority groups: $m_1, m_2, \ldots, m_k, k > 2.$

⁷For the Atkinson index, it is the case because its parameter b is set to 0.5.

⁸"we artificially restricted attention to the two-group cases. Indices were calculated as if Anglos and the minority in question were the only two groups present. For example, in measuring the degree of segregation between Hispanics and Anglos, the total population was assumed to equal the sum of these two groups", Massey and Denton (1988), page 299.

- 3. *Concentration* "refers to the relative amount of physical space occupied by a minority group in the urban environment" (MD, page 289);
- 4. *Centralization* "is the degree to which a group is spatially located near the center of an urban area" (MD, page 291);⁹
- 5. *Clustering*, the degree of spatial clustering exhibited by a minority group: "the extent to which areal units inhabited by minority members adjoin one another, or cluster, in space" (MD, page 293).

Compared to the exhaustive perspective of MD, we focus on the first two features of segregation: evenness (or, instead, randomness as explained above) and exposure. The other dimensions emphasized by MD are geographical. They are pertinent in the study of residential segregation in a given city but appear less applicable when estimating an index of segregation at the level of an entire country, what we do here for metropolitan France.

The most common indices associated with evenness/randomness are the Duncan, the Theil, the Atkinson, and the Gini indices. The classical index for exposure is the Coworker index.¹⁰ Those indices differ as regards the composition-invariance property (James and Taeuber (1985), that is, whether the index is sensitive to the relative proportions of the groups m_1 and M: "Although indices of exposure and evenness tend to be correlated empirically they are conceptually distinct because the former depend on the relative size of the groups being compared, while the latter do not. [...] Exposure indices take explicit account of the relative size of minority and majority groups in determining the degree of residential segregation between them" (MD, page 287).

In the following analyses, we estimate both types of indices. These estimations enable us to compare the evenness/randomness dimension with the exposure dimension in the terminology of Massey and Denton (1988). We discuss that point in the presentation of our results. Henceforth, we reserve the term *dimension* for the different types of demographic or socioeconomic variables used to define minority and majority groups. Finally, remark that this distinction between composition-invariant and composition-variant indices is also known as the distinction between the *intensity* ("intensité") of segregation for the Duncan (or another composition-invariant index) and the *scale/extent* ("ampleur") of segregation for the Coworker index (see Shon and Verdugo (2014)).

Literature

This article connects to a broad literature on the measure of segregation. The articles that address the small-unit bias have already been mentioned. In addition to those already cited, other methodological articles include: Duncan and Duncan (1955) that defines the dissimilarity or Duncan index and discusses its links with the segregation curve; Reardon and Firebaugh (2002) for the discussion of multi-group settings; Frankel and Volij (2011) for axiomatic properties of proportion-based indices; Åslund and Skans (2009) for conditional indices.

⁹The notion appears more relevant in the U.S. context of ethnic residential segregation, as the more deprived population often concentrates near the center of the city. It would rather be the opposite in a European context.

¹⁰Remark that we reuse here the terminology of MD, but the exposure dimension is also subject to the distinction between realized proportions (evenness benchmark as discussed above) versus underlying probabilities (randomness benchmark). In this article, we always consider the randomness benchmark.

Regarding the application, several articles are closely linked to our work. Maurin (2004) also uses the LFS data to obtain concentration measures, but distinct from the classical segregation indices due to the small-unit bias, for social status and ethnicity dimensions. Rathelot (2012) also uses the LFS data and a probability-based approach to estimate segregation indices in the ethnicity dimension over the period 2005-2008. Shon and Verdugo (2014) studies the residential segregation of immigrants in France. It does not rely on LFS data but on censuses. Consequently, the role of the units is not played by the 20 to 40 adjacent housing clusters of the LFS but by residential IRIS ("Ilots Regroupés pour l'Information Statistique", aggregated units for statistical information), namely infra-municipal areas of around 1,800 to 5,000 inhabitants. Although IRIS are defined to be homogeneous in terms of living environment with boundaries based on the major dividing lines provided by the urban fabric (main roads, railways, bodies of water, etc.),¹¹ the corresponding neighborhood areas are far coarser than the clusters of the LFS. Although the latter requires tackling the small-unit bias, we believe it is worth the cost to study residential segregation because the LFS clusters provide arguably natural neighborhoods of around fifty adults living in adjacent housing, for each year (even each quarter since 2003), and with rich individual covariates, which allows exploring multiple dimensions of segregation.

Lastly, by considering the ethnicity dimension, this work also relates to a strand of the literature studying the job market or residential conditions of immigrants or ethnic groups proxied by parents' nationalities or places of birth, such as Aeberhardt et al. (2010) and Meurs et al. (2006).

Main findings

The results suggest two main conclusions. First, whatever the dimension under study, (i) labor market positions, (ii) social/occupational categories, (iii) levels of education, (iv) ethnicity, (v) origins (immigrants), or (vi) nationalities, the estimated indices do not reveal significant evolution over time: *within each dimension*, the magnitude of residential segregation has remained globally constant for the past decades.

Second, the magnitude of segregation differs *across the different dimensions* according to this decreasing ranking: (vi) nationalities, then (v) origins (immigrants) and (iv) ethnicity, followed by (iii) education levels and (ii) social/occupational categories, finally (i) labor market positions with the lowest magnitude of segregation.

A conditional analysis, separating neighborhoods that belong to urban areas of 200,000 inhabitants or more from neighborhoods belonging to smaller urban areas, overall confirms these findings for the aggregate conditional indices. Nonetheless, this conditional analysis also suggests that, in some dimensions, the temporal evolution of segregation differs between large cities (200,000 inhabitants or more) and smaller cities or rural areas.

2.2 Data and dimensions

2.2.1 Labor Force Survey and units/neighborhoods

The data used is the French Labor Force Survey (LFS) from 1968 to 2019. Before 2003, the survey was annual; since 2003, data has been available each quarter. For the period

¹¹See INSEE documentation https://www.insee.fr/en/metadonnees/definition/c1523 for further details.

2003-2019, we restrict to the first quarter to present the evolution with yearly frequency. The survey field is constituted by principal (by opposition to secondary residences) ordinary housing within which are surveyed every individual who is at least 15 years old.

The LFS is one of the leading surveys conducted by the French National Institute of Statistics and Economic Studies (INSEE), and the sample sizes are large. For instance, regarding the rotating panel, between 2003 and 2009, on average, 54,000 ordinary housings are surveyed each quarter, 44,000 are in the survey field (i.e., are principal ordinary housing), 36,000 answer the survey, which results in interviewing about 70,000 individuals of more than 15 years old. In 2009, the sample size was increased gradually by 50%. In 2013, 57,000 ordinary principal housings counting for about 108,000 individuals older than 15 years old, were interviewed.

In the analyses, we restrict to adult individuals (at least 18 years old) with non-missing individual variables used to define the minority and majority groups. Also, we restrict units containing at least two distinct housing, which is the case for virtually all units unless rare exceptions.

In all analyses, a unit or neighborhood is defined as a cluster of the areal sampling scheme of the LFS. It is a geographic area of about 30 adjacent ordinary housing, representing 48 adults living in the neighborhood, on average over the period 1968-2019.

Depending on population densities, notably rural versus urban environments, two clusters can have quite different physical delimitations. In rural areas, it could correspond to an entire hamlet or different "lieux-dits" (small localities composed of a few housings) that can be separated by several hundreds of meters or a few kilometers. In residential suburbs of moderate density or small to medium rural towns, a cluster is typically a neighborhood of side-by-side houses (see Figure 2.1, a picture from Google Maps based on a real example given in the survey instructions of the 2008 LFS). In large cities, a cluster can be one or two buildings of a street or even, for large buildings, one or two floors of flats.

Those differences in the physical size of clusters can entail different subjective meanings of these clusters for individuals who live in. For instance, in a large city, the inhabitants of the three buildings next to mine that belongs to my cluster may not mean more to me than other inhabitants of buildings further away who are not in my cluster or other city inhabitants I see taking public transportation. On the contrary, in medium towns, clusters like the one of Figure 2.1 have probably an immediate sense of neighborhood for individuals who live in. Despite those differences, the LFS areal sampling clusters appear as relevant units to study residential segregation. For units composed of one or two buildings in dense areas, we may consider that they are somewhat too small to represent the neighborhood of an individual fairly; nonetheless, the unit's composition conveys information on the composition of the immediate neighborhood.

2.2.2 Dimensions of segregation and individual covariates

This section describes the different dimensions of segregation that we study, namely the individual covariates used to define minority and majority groups.

(i) Labor market positions

The first dimension we study is the labor market positions. This dimension relies on an individual covariate that indicates the labor market position or status with the International



Figure 2.1: Example of a cluster (2008 LFS) in an area of moderate population density.

Labor Organization classification: active (that is, member of the labor force) employed, active jobseeker, or inactive (that is, outside of the labor force). That variable is available through the entire period 1968-2019 and is rarely missing. The idea of this dimension is to assess the possible residential segregation of jobseekers.

We consider two specifications for this dimension d = (i) labor¹²

- (a) minority group m = unempl. versus majority group M = empl., either we restrict to individuals active in the labor market, so that the minority group is made of unemployed people and the majority group of employed individuals;
- (b) minority group m = unempl. versus majority group M = empl. or inact., or we encompass the entire population with a minority group that remains made of unemployed individuals while the majority group gathers active employed individuals and inactive individuals, out of the labor force.

Remark that, in specification (a), as explained in Section 2.1, §"Dimensions of segregation" with the different groups m_1 , m_2 and M, the individuals out of the labor market are dropped from the analysis.

(ii) Social/occupational categories

The second dimension relates to the French classification of "professions et catégories socio-professionnelles" (PCS/CSP). Since 1954, it classifies individuals according to their professional category or situation, considering several criteria: profession, economic activity, qualification, hierarchical position, and status. A significant recast of the classification was done in 1982. Despite this, it is possible to keep track of the particular category we consider for defining the minority group in this dimension: managerial and intellectual professional occupations ("cadre et professions intellectuelles supérieures" in French), such as liberal

¹²The denominations written in typewriter font for dimension d, minority group m and majority group M are used as abbreviations in the figures presenting the results.

professions (physicians, lawyers, etc.), civil service managers, professors and scientists, professions in information and communication technologies or arts, administrative and commercial managers, engineers, etc. The variable is available throughout the 1968-2019 period.

We consider two specifications for this dimension d = (ii) occup.

- (a) minority group m = manag. versus majority group M = other act., either restricting to the individuals currently active in the labor market, where the minority group is made of active individuals occupying a managerial or intellectual professions and the majority group of the other individuals active in the job market;
- (b) minority group m = act. or former manag. vs majority group M = other, or also including in the analysis former managerial and intellectual professions, so that the minority group is composed of active and former managerial and intellectual professional individuals while the majority group is made of the rest of the population with known current or former (for retired individuals) occupational categories.

(iii) Education levels

This dimension is defined by the highest level of education, or diploma, attained by individuals. Compared to all other dimensions where we include in the analysis adult individuals (at least 18 years old), for this dimension, we restrict to individuals at least 23 years old because the preceding years are typically those of the initial formation. The variable is available from 1970, except for the year 1973.

We consider a single specification for this dimension d = (iii) educ.

(a) minority group m = bachelor or + versus majority group M = other. The minority group m is made of individuals with at least a bachelor's degree, that is, in French equivalent, a "Licence" (a three-year degree in the LMD European system of 2004 – Bachelor/Licence, three years; Master, over five years; Doctorate, over eight years). The majority group M is made of all other individuals (aged at least 23), whose highest achieved education level is lower than a bachelor's degree.

(iv) Ethnicity

This dimension is defined from a proxy of the ethnicity of individuals. The use of a proxy is necessary because French law prohibits collecting race or ethnicity statistics. For a given individual, the usual proxy is based on his or her parents' nationality or place of birth at the parents' birth. In the LFS, the modalities for those variables (nationality or place of birth) of parents are the following:

- 1 France,
- 3 Northern Europe,
- 4 Southern Europe,
- 5 Eastern Europe,
- 6 North Africa (Algeria, Morocco, Tunisia),
- 7 Rest of Africa (Africa except the three countries/nationalities of modality 6)
- 8 Middle-east
- 9 South-East Asia, which corresponds to Laos, Vietnam, and Cambodia (former French colonies)
- 10 Rest of the world.

We indicate the modality as encoded in the LFS documentation as we will refer to that encoding when defining minority and majority groups.

Different specifications can proxy ethnicity depending on the choice of nationality or place of birth to determine the modality of an individual's parents and the choice of the father or the mother as the dominant parent. Overall, they yield very similar results regarding the proxied ethnicity of individuals. Therefore, we focus on one of the 2×2 possibilities, which is to take the mother (instead of the father) as the dominant parent in a sense detailed below and the place of birth (instead of the nationality) at the birth of the parents.

An individual proxied ethnicity takes values in the above-mentioned modalities (1 to 10, without 2) and is defined in the following way:

- whenever the two parents have modality 1, the proxied ethnicity is 1;
- when one of the two parents (be it the father or the mother) has modality 1 and the other parent a modality $x \neq 1$, the proxied ethnicity is x;
- finally, if the two parents have modality x (father) and y (mother) different than 1, if x = y, the proxied ethnicity is x = y; if $x \neq y$, the proxied ethnicity is the one of the dominant parent, here chosen to be the mother, hence y.

The variables required to obtain the proxied ethnicity have only been available since 2005.

We consider three specifications for this dimension d = (iv) ethni.

- (a) minority group m = a.m.t. versus majority group M = french, where minority individuals are those whose proxied ethnicity has modality 6 (Algeria, Morocco, or Tunisia, abbreviated to a.m.t.) and majority individuals those with modality 1 (the two parents of the individual are born in France);
- (b) minority group m = african versus majority group M = french, minority individuals with proxied ethnicity equal to 6 or 7 (at least one of the parents of the individual is born in Africa) and majority group as in (a);
- (c) minority group m = non-eur. versus majority group M = european, minority individuals with proxied ethnicity equal to 6, 7, 8, 9, or 10 and while the proxied ethnicity of majority individuals belong to the modalities 1, 3, 4 or 5.

(vi) Nationalities

In the LFS, the modalities encoding nationalities vary over time with more or less precision. This explains the following specification choices and the lack of availability of some of them over the entire period 1968-2019. For easier comparison, we express the modalities of the nationalities with the modalities of proxied ethnicity.

We consider four specifications (the first three are analogous to dimension (iv) with nationalities instead of proxied ethnicity) for this dimension d = (vi) nat.

(a) minority group m = a.m.t. versus majority group M = french, individuals of Algerian, Moroccan or Tunisian nationalities versus individuals of French nationalities, the variables and related modalities required for that construction are available for the full period 1968-2019;¹³

¹³This is probably explained by the colonial links between France and these countries.

- (b) minority group m = african versus majority group M = french, African nationalities versus French nationality, the related modalities for this specification are only available from 1975;
- (c) minority group m = non-eur. versus majority group M = european, non-European versus European nationalities, that specification is available from 1990 only;
- (d) minority group m = non-fr. versus majority group M = french, French nationality as the majority group contrasted to all other foreign nationalities in the minority group, that specification is available for the entire period 1968-2019.

(v) Origins (immigrants)

This dimension is connected with the previous two dimensions. Indeed, for a given reference, here France, an immigrant is defined as an individual who resides in France and who is born abroad (outside of France) with a foreign nationality at birth (not French). Note that the criterion is defined once at the birth of the individual. In particular, an individual can be an immigrant but have acquired French nationality by naturalization.

We consider two specifications for this dimension d = (v) orig.

- (a) minority group m = immigrant versus majority group M = other, where an individual belongs to the minority group if he or she is an immigrant in the exact definition of being born as a foreigner abroad while the majority group is made of all other adults; the variables needed for this construction have been available since 2003 only;
- (b) minority group m = foreign immi. versus majority group M = other, where an individual belongs to the minority group if he or she is an immigrant *and* does not have the French nationality when surveyed (thus, it is only a proper subset of the set of immigrants of specification (a)) and the majority group is made of all other adults; the required covariates to construct those groups are available in the data between 1982 and 2019.

2.3 Methodology and robustness

2.3.1 Parameters of interest, estimation, and inference

This subsection briefly recalls the statistical model, the index of interest, and how estimation and inference are performed. We refer to the original articles by Rathelot (2012) and D'Haultfœuille and Rathelot (2017) as well as the Stata article D'Haultfœuille et al. (2021) that presents the package segregsmall used in this application for further details.

The population of interest is split into a minority group m and a majority group M according to some dimension and specification (see Section 2.2.2). We assume to observe an independent and identically distributed (i.i.d.) sample $(X_i, K_i)_{i=1,...,n}$ of units, with X_i the number of minority individuals living in unit i and K_i the total number of individuals (minority m plus majority M) living in neighborhood i.

We start by reasoning conditional on K = k.¹⁴ A probability-based segregation index θ is then a functional of the cumulative distribution function F_p of p and of $m_{01} = \mathbb{E}[p]$, that

¹⁴As we assume *i.i.d.* sampling of units/neighborhoods, we drop the subscript *i* to lighten notation: $(X_i, K_i, p_i)_{i=1,...,n} \stackrel{\text{i.i.d.}}{\sim} (X, K, p).$

is $\theta = g(F_p, m_{01})$. As explained in Section 2.1, the index will be minimal equal to 0 when p is constant (Dirac mass) and maximal equal to 1 when p follows a Bernoulli distribution.

The function g depends on the mapping used to quantify the variation of the probabilities p_i . We consider five classical segregation indices

$$\begin{aligned}
\theta^{\rm D} &= \frac{\int |u - m_{01}| \, \mathrm{d}F_p(u)}{2m_{01}(1 - m_{01})} & (\text{Duncan}), \\
\theta^{\rm T} &= 1 - \frac{\int \{u \ln(u) + (1 - u) \ln(1 - u)\} \, \mathrm{d}F_p(u)}{m_{01} \ln(m_{01}) + (1 - m_{01}) \ln(1 - m_{01})} & (\text{Theil}), \\
\theta^{\rm A(b)} &= 1 - \frac{m_{01}^{\frac{-b}{1-b}}}{1 - m_{01}} \left\{ \int (1 - u)^{1-b} u^b \, \mathrm{d}F_p(u) \right\}^{\frac{1}{1-b}} & (\text{Atkinson, with } b = 0.5) \\
\theta^{\rm CW} &= \frac{\int (u - m_{01})^2 \, \mathrm{d}F_p(u)}{m_{01}(1 - m_{01})} & (\text{Coworker}), \\
\theta^{\rm G} &= \frac{1 - m_{01} - \int F_p^2(u) \, \mathrm{d}u}{m_{01}(1 - m_{01})} & (\text{Gini}).
\end{aligned}$$

For the moment, K is constant or fixed in a conditional analysis. In practice, the unit sizes vary. To deal with random K that takes values in \mathcal{K} , the segregation index θ is defined as a weighted average of indices conditional on K = k, denoted $\theta^k = g(F_p^k, m_{01}^k)$ with F_p^k the cdf of p conditional on K = k, and $m_{01}^k = \mathbb{E}[p \mid K = k]$; hence, θ depends on the joint distribution of p and K. Whether we study segregation at the unit-level or at the individual-level matters for the weights used. The unit-level index θ_u satisfies

$$\theta_u = \sum_{k \in \mathcal{K}} \mathbb{P}(K = k) \, \theta^k$$

whereas the individual-level segregation index θ_i is defined by

$$\theta_i = \sum_{k \in \mathcal{K}} \frac{k \mathbb{P}(K=k)}{\mathbb{E}[K]} \theta^k.$$

We use either the non-parametric (np) approach of D'Haultfœuille and Rathelot (2017) or the Beta (*beta*) parametric approach of Rathelot (2012) as implemented in the segregsmall Stata command to estimate θ_u and θ_i for each of the five classical segregation indices.¹⁵ Confidence interval on θ_u and θ_i are obtained by bootstrap using 50 repetitions.¹⁶ The parametric method *beta* assumes a Beta distribution for F_p and yields point-identification and usual point-estimators for the target parameters θ_u and θ_i obtained by maximum likelihood. On the contrary, in general, θ_u and θ_i are only partially identified in the non-parametric np approach, which does not restrict F_p . In concrete terms, it means that, instead of a usual point-estimate for the segregation index θ_u or θ_i , the np method reports an estimated identified interval for the index.¹⁷ However, it happens that for sufficiently large K (in practice, K of around a dozen is often enough),

¹⁵Remark that the Gini index cannot be estimated through the non-parametric as the related map g does not satisfy a required condition to implement the np approach.

¹⁶This relatively limited number was constrained by computational time requirements. Nonetheless, for some specifications, the comparisons of confidence intervals obtained with 50 or 100 bootstrap repetitions suggest the changes from a larger number of bootstrap repetitions are negligible compared to the statistical uncertainty captured by the confidence interval.

¹⁷If you are not familiar with partial identification, think about the following simple example. Imagine you want to estimate the unemployment rate in France; actually, this is the primary goal of the LFS

the index is, in fact, point-identified. In this case, the lower and upper estimated bounds coincide, and the np method outputs a point-estimate for the index $\theta = \theta_u$ or θ_i . Since, on average, over each year, the 5%-quantile of K is equal to 11, in our application, most of the time, the non-parametric approach yields point-estimators, although it is not always the case, depending on the exact specification and the year.

2.3.2 Specification choices and robustness

The remaining choices needed to obtain estimations of segregation indices are the following.

Method The method used for estimation: either the non-parametric approach of D'Haultfœuille and Rathelot (2017) that leaves the distribution F_p unrestricted (np) or the parametric approach of Rathelot (2012) that assumes a Beta distribution for F_p (*beta*). Remark that for the Gini index, only the *beta* method is available.

Weight The weights used for defining the segregation index when aggregating over random K: either unit-level weights for the parameter θ_u (unit) or individual-level weights for θ_i (ind.). Since we are interested in a global measure of residential segregation estimated from units with varying sizes, individual-level weights that give more weight to the contribution of larger neighborhoods in the aggregated index appear to be preferred in principle. Nonetheless, we might wonder at this stage what the impact of that choice is (see robustness checks below).

Assumption on (K, p) Estimation and inference can be made under two settings regarding the dependence between the variables K and p: without assumption, leaving any possible dependencies between the two variables, or assuming independence. A priori, it is more careful to be agnostic about the joint distribution of (K, p). Nonetheless, independence allows for pooled estimation, improving identification, and reducing computational requirements. Besides, under this assumption, we have $\theta_u = \theta_i$, and the chosen weights are thus irrelevant. Sampling scheme concerns partly determine the sizes K of units. For instance, the clusters are constructed for easier exploration and finding of newly constructed housing to ease the work of INSEE interviewers. Also, the number of individuals included in the analysis depends on the number of observed individuals, thus partly depending on the experience and perseverance of interviewers to obtain answers. Such factors related to the sampling procedure seem a priori unrelated to possible segregation, which provides arguments for assuming $K \perp p$.

Sample of units The sample of units/neighborhoods included in the analysis. As discussed in D'Haultfœuille et al. (2021) (Section 3.5), the complexity and computational time of the np method as implemented in the command segregsmall are primarily

data used in this article. You have drawn randomly and independently 100 individuals representative of your target population. Thanks to that sampling scheme, you can assume the observations as *i.i.d.*. However, only 90 individuals answered your questionnaire. Among those 90 individuals, 14 are jobseekers. Let assume that, whatever the sample size and your efforts, there will always be a 10% fraction of non-respondents. In such a case, without additional assumptions such as missing-at-random selection or imputation of missing values, the unemployment rate is only partially identified, you cannot state more than the identified interval [4, 24] = [14 - (100 - 90), 14 + (100 - 90)] whose bounds are obtained if either all the non-respondents are employed (lower bound) or they are all jobseekers (upper bound).

influenced by the magnitude of K: the higher K and the larger $|\mathcal{K}|$, the more demanding the optimization problem. To study the evolution of segregation over time and compare its magnitude across dimensions, we make the analysis for each year (52 years for the whole period 1968-2019) and several dimensions and specifications. In addition, each time, we compute confidence intervals from bootstrap. Consequently, we cannot afford an optimization that requires too much computational time with the np method. In practice, the computational requirements are too demanding for our desired application for Ktaking each integer value between, say, K = 10 and $K \approx 50$ or 60.

One possibility to deal with that computational burden would be to favor the parametric *beta* approach that is faster, notably under the assumption $K \perp \rho$.

Another is to restrict the units included in the analysis to the units with size $K \leq \overline{K}$ with $\overline{K} := 50$. At first sight, this may seem like a strong and arbitrary restriction. However, given the distribution of K, such restriction to small neighborhoods encompasses a (vast) majority of all neighborhoods. Between 2003 and 2019, pooling all years together, the proportion of units with a number of adults at most equal to 50 is 97%. Before, due to the larger units in the first part of the period, the proportion is inferior: 67% of all units between 1990 and 2002 but only 34% between 1968 and 1989. However, note that these figures are computed using the raw number of adults in the units. In practice, as explained in the construction of the different dimensions of the analysis in Section 2.2.2, the total number K of minority m plus majority M individuals per unit included in the analysis is lower: some, but relatively few, individuals have missing data for the required covariates; more importantly, in several specifications, we restrict the population to a minority group $m = m_1$ and a majority group M and drop other groups m_2, \ldots, m_k as explained in Section 2.1. As a result, the restriction to units with $K \leq \overline{K} = 50$ always comprise the majority (often an overwhelming, especially in recent years) of the units.

Nonetheless, to mitigate that concern, we also consider another type of sample selection. The idea is the following. Units with $K \leq \overline{K}$ are included as such in the analysis. For units with $K > \overline{K}$, we randomly draw \overline{K} among the K individuals of that unit; this form a new (sub-) unit with $K = \overline{K}$ individuals, which is included in the analysis. That idea might seem surprising first as we introduce additional variability in the composition of such units: we may randomly pick, say, more minority individuals in the subsample of selected individuals than their proportion in the entire unit. However, first, there is no reason this happens systematically across units; second, by definition of the randomness benchmark, we are not interested in the realized proportion X/K but the underlying unobserved probability p; it is simply that the latter will be retrieved from a smaller sample of individuals, as it is already the case in smaller units. In fact, this type of subsampling corresponds to the particular case of survey data discussed in D'Haultfœuille et al. (2021) (Section 2.2, §"Assuming independence between K and ρ ") and in D'Haultfœuille and Rathelot (2017) (Section D.4) In other words, we do lose some statistical power with that selection but in exchange for a lesser computational burden.

Finally, to assess the influence of such a selection of units (first case) or of the subsampling of individuals within the large units with size exceeding \overline{K} (second case), we also compare the obtained estimates from the np method with the estimate computed on the entire set of units in the particular case of the Coworker index. Indeed, conditional on K, thanks to the specific map g defining the Coworker index, the index only depends on the first two moments of the distribution of $p \mid K$ and can thus be easily computed even for large and rich \mathcal{K} .

Robustness checks

As just explained, the estimations of segregation indices depend on four specification choices about

- the method used (*beta* or np);
- the weights used to compute the aggregate segregation index $(\theta_u \text{ or } \theta_i)$;
- the assumption made on the joint distribution of (K, p) (assumed independent or unrestricted);
- the sample of units/neighborhoods included in the analysis (all, but not available in practice due to computational burden, or a restriction to the small units $K \leq \overline{K} = 50$, or a subsampling of individuals for the large units $K > \overline{K} = 50$).

In this subsection, we summarize the results of several robustness checks implemented to assess the impact of those four choices. For each of them, the conclusion is that the choices made have no impact on the qualitative conclusions of the study of residential segregation and limited impact on the quantitative findings. Different choices slightly change the quantitative values of the estimates, but the differences are rather small, and confidence intervals for the different choices mostly overlap so that we cannot reject those choices lead to significantly different values of segregation indices.

The rest of this subsection presents three figures to illustrate this point. The absence of real impact of the methodological choices in the conclusions hold similarly across the different dimensions of segregation studied and the distinct indices (Duncan, Theil, Atkinson, Coworker, Gini) as regards method (np or beta) weights (unit or indiv) and assumptions on (K, p) (none or assumed independence). For this reason and for the sake of conciseness, we only present some of them.

For the three different samples of units included in the analysis: small units with size at most \overline{K} (small ($K \leq 50$)), all units with sampled individuals for the large ones (all (smpl. if K > 50)), all units as such, without subsampling of individuals (all), there is the caveat that the comparison is only made for the Coworker index. For that index, the three samples of units give very similar results (Figure 2.3). It is not immediate to find reasons why it would be different for the other indices. Although not a proof, we are thus relatively confident about the robustness of our analysis relative to this choice of sample selection for the other segregation indices. We acknowledge, however, that more direct robustness checks would be appreciable; we leave them for future work.

Figure 2.2 illustrates the limited impact of the choices relative to weights, assumptions between K and p, and the sample of units included in the analysis, at least for the choice of small units and the choice of all units with a subsampling of individuals in large units. The estimates of θ_u (unit-level weights) seem slightly above the estimates of θ_i (individual-level weights), but moderately. Remark that under the assumption $K \perp p$, $\theta_i = \theta_u$ and the related estimators also coincide. This explains why the solid red and green lines (*ind.*) are not visible: they are hidden by the solid blue and purple ones (*unit*) since the estimates are equal.

In addition, the difference between the two samples of units decreases over time. This is because, as already mentioned, the size K decreases between 1968 and 2019. Consequently, they are fewer and fewer units with more than \overline{K} individuals and, thus, the two choices for the sample used get closer to each other.

That point is also visible in Figure 2.3, which focuses on the impact of the choice of the sample of units included in the analysis. As already stressed, it is performed for Figure 2.2: Dimension: (ii) social/occupation categories with specification (a) restriction to active individuals. Index: Duncan. For the np method, the figure compares the remaining three methodological choices: weights, assumptions between K and p, and sample of units included in the analysis (only small ($K \leq 50$) and all (smpl. if K > 50)).



the Coworker index to compare all three possible choices. The comparison is made for individual weights, np method and without independence assumption between K and p. It illustrates the limited impact of the sample selection, notably in recent years since the switch to the rotative and continuous LFS in 2003.

For the Duncan index, Figure 2.4 compares the non-parametric np and parametric beta approaches, as well as weights in the case of the absence of independence assumption between K and p and for the sample of all units with subsampling of individuals for the large ones. For given weights, the target parameter of both methods beta and np is the same (either θ_u or θ_i). De facto, the estimates are very close: a Beta assumption for F_p leads to almost the same estimated segregation index as a non-parametric method. This is reassuring, in particular for the Gini index. which cannot be computed with the npmethod as it fails to satisfy some regularity assumption (D'Haultfœuille and Rathelot (2017), Assumption 2.1). If available, np method is in principle preferable given its flexibility despite that it is more expensive in terms of numerical computations than the parametric approach. Although not direct proof, the proximity of both methods for the Duncan (the same holds for the Theil, the Atkinson, and the Coworker) brings confidence to a correct estimation of the Gini index by the *beta* method. Finally, Figure 2.4 also illustrates the possibility of partial-identification of the index with the non-parametric approach. In this case, np methods give an estimated identified interval for the index. Such an interval can be seen, for instance, in 2010 for θ_i (individual weights).

2.4 Results

This section shows several graphs of annually estimated segregation indices to study temporal evolution and compare the magnitude of segregation across different dimensions. Figure 2.3: Dimension: (i) labor market positions with specification (a) restriction to active individuals. Index: Coworker. For the absence of independence assumption between K and p, np method, and individual-level weights, the figure compares the three possibilities of the remaining methodological choice as regards the sample of units included in the analysis: *small* ($K \leq 50$), all (*smpl. if* K > 50) and all.



Figure 2.4: Dimension: (ii) social/occupation categories with specification (a) restriction to active individuals. Index: Duncan. For the absence of independence assumption between K and p, the choice of all (smpl. if K > 50) for the units included in the analysis, the figure compares the remaining two methodological choices: weights (unit or ind.) and methods (beta or np).



2.4.1 Unconditional analysis

Combined with previous explanations, the graphs' legends and subtitles should be self-sufficient to understand:¹⁸

- 1. the specific segregation index considered (see Section 2.3.1);
- 2. the methodological choices made regarding method, weights, assumption on (K, p) and sample (see Section 2.3.2);
- 3. the dimension of segregation and the specifications of minority and majority groups (see Section 2.2.2).

For a given index (1.), methodological choices (2.), and definitions of groups (3.), estimation and inference on the segregation index is performed separately for each year. Thus, one point on the following figures corresponds to one estimate (or estimated lower/upper end of the identification interval for np method in case of partial-identification) of the index; the solid lines connect them. If appropriate, the y-title of the graph indicates the asymptotic level of the confidence intervals (mostly 95%), which are displayed in dashed lines. Note that those confidence intervals are only pointwise or cross-sectional intervals, valid for each year (instead of longitudinal interval tubes). Finally, the y-axis scale and grid adjust to each graph; comparisons across different figures should thus be made with some caution.

The subtitle of each graph specifies the index (1.) and the methodological choices (2.). Remember that the possible indices are the Duncan, the Theil, the Atkinson, the Gini, or the Coworker. They are all normalized between 0 and 1, with higher values indicating larger magnitudes of segregation.

The four first indices are composition-invariant: they quantify the concentration of minority individuals in some neighborhoods more than in others irrespective of the fractions of the minority and majority groups in the population.

Overall, the estimation of those four indices yields the same qualitative conclusions regarding the temporal evolution and the comparison of the magnitude of segregation across dimensions. Nonetheless, the quantitative values of the estimates differ as the indices rely on distinct mapping g to synthesize the variations of the probabilities into a scalar index. In practice, the estimated Duncan and the Atkinson indices have similar values, ranging in our application over almost the entire possible range, from around 0.1 to 0.9. The Gini index shows similar numerical values, perhaps slightly larger than the Duncan and the Atkinson, with estimates ranging from 0.3 to 0.9. In contrast, the quantitative values of the Theil index are substantially smaller, around half those of the Duncan and Atkinson indices.

Again, despite those changes, the conclusions regarding temporal evolution and comparisons across dimensions are overall consistent between these four indices. That is why we do not always report here the four of them for the sake of conciseness.¹⁹

On the other hand, the Coworker index is defined as a measure of exposure in the classification of Massey and Denton (1988): it is sensible to the proportions of the minority and majority groups in the population. All else equal, the closer those proportions to one half, the larger the Coworker index. Intuitively, this relates to the interpretation of the Coworker index as the mean exposure of minority individuals to other minority individuals (that is, the probability for a minority individual to reside in a neighborhood with at least another member of the minority group).

¹⁸That is why the captions of the figures, if given, are rather comments or remarks instead of explanations. ¹⁹Some additional forward are presented in Appendix 2. Λ

¹⁹Some additional figures are presented in Appendix 2.A.

Consequently, because it does not measure the same aspect of the segregation phenomenon, the Coworker index might show different temporal evolution or ranking across dimensions than the composition-invariant indices (Duncan, Atkinson, Theil, and Gini). It happens to be the case in this application to residential segregation in France. In particular, in some specifications, the Coworker displays an increasing trend over time. At first glance, this contradicts the overall stability of segregation over time within each dimension revealed by the other indices. Nonetheless, it is not the case because, for a given level of the concentration of minority individuals in some units only (that is, a given magnitude of segregation as quantified by a composition-invariant index), the Coworker index changes if the population proportions of the two groups evolve. For instance, in the labor market position dimension, the increase in the Coworker index over time is essentially due to the rise in unemployment over the period 1968-2019.

There have been long-standing debates about the desirability for a segregation index to be sensitive to the proportions of minority and majority groups in the population. It is the case for the Coworker index, but not for other classical indices we study here (Duncan, Atkinson, Theil, and Gini). In this article, we tend to consider that the magnitude of segregation should relate to the uniform or concentrated distribution of minority individuals across neighborhoods, irrespective of the size of the minority group in the population. Consequently, our main findings rely on the estimated composition-invariant indices, although we also study and present our results for the Coworker index for completeness and comparisons.

Regarding methodological choices (2.), Section 2.3.2 explains that they have a limited impact on the results. Consequently, we primarily focus on our preferred specification: non-parametric np method, no assumption as regards possible links between K and p(those two choices ensure the maximal flexibility), individual-level weights (the index of interest is θ_i), and all units included in the analysis with subsampling of individuals in the large ones to alleviate numerical burden (if their sizes exceed $\overline{K} = 50$). Appendix 2.A present some figures for other methodological choices for comparison. The choices are indicated in the subtitle of each figure, with some (minimal) abbreviations introduced in the first paragraphs of Section 2.3.2.

Finally, the legends of each figure indicate the dimension of segregation and the precise specifications of minority and majority groups (3.) for which segregation estimates are plotted.²⁰ Each legend item contains three elements separated by a slash. The first one is the dimension d of segregation, that is, essentially, the type of demographic or socioeconomic variables used to construct group membership. The second and third elements respectively indicate the minority group m and the majority group M that are studied. The abbreviations used are presented in Section 2.2.2 when defining the different dimensions and minority/majority specifications.

We recall that the union of m of M does not necessarily cover the entire conceivable population. If not, the individuals who belong neither to m nor to M are dropped from the analysis that restricts to two-group settings with minority m and majority M (see Section 2.1, §"Dimensions of segregation").

If a single figure was to sum up the unconditional analysis, we may elect Figure 2.5. It shows the estimated Duncan indices for each year (solid lines) and associated 95%

²⁰The color associated with a given specification may change across figures depending on the different dimensions that are represented; comparisons across figures should be made with care.

confidence intervals (dashed lines) for the six different dimensions and our preferred specification regarding methodological choices.

It illustrates the two main findings that synthesized our unconditional analysis. First, within each dimension, the evolution of the magnitude of segregation has remained overall constant for the past decades. Second, across the different dimensions, the magnitude of segregation differ sharply with the following decreasing ranking, from highest to lowest magnitude of segregation,

- (vi) nationalities, with a Duncan index around 0.70;
- followed by (v) origins (immigrants) and (iv) proxied ethnicity, with similar magnitude displaying a Duncan index about 0.50;
- then (iii) education levels and (ii) social/occupational categories with a close magnitude of 0.40 for the Duncan index;
- finally (i) labor market positions with a Duncan index around 0.25.

Figure 2.5: illustrates the two main findings stated in Section 2.1, §"Main findings."



As discussed above, the quantitative values of segregation indices are partly arbitrary as they depend on the choice of the mapping g to obtain a normalized index between 0 and 1. For instance, in contrast to Figure 2.5 (Duncan index), Figure 2.6 shows the estimates of the Theil index with the same methodological specifications and for the same dimensions (but without confidence interval for easier readability). Despite that the quantitative values of the indices differ (with the Theil index being smaller), the qualitative conclusions are unchanged. That is why, arguably, it is, above all, interesting to compare indices over time or across dimensions rather than to focus on their exact quantitative values.²¹

That being said, the quantitative value of the Duncan index can be interpreted, and that interpretation can furnish a concrete intuition on the magnitude of segregation. Indeed, the value of the Duncan index corresponds to the proportion of minority individuals who would

²¹Remark that, although its quantitative value in itself might not be the most interesting part of a segregation index, valid estimation, notably addressing the small-unit bias, is nonetheless crucial to be able to make reliable comparisons of indices.



Figure 2.6: same as Figure 2.5 for the Theil instead of the Duncan index.

Index: Theil – Units: all (smpl. if K > 50), weights: individual, method: np, (K, p): no assumption

have to change units (namely, move from their neighborhoods for residential segregation), that is, quit the units where the minority group are over-represented to join units where they are under-represented so as to achieve a uniform distribution across units of minority individuals (Duncan and Duncan, 1955). That is why the Duncan or dissimilarity index is sometimes named displacement index. Note that such an interpretation makes sense in an evenness benchmark as opposed to the randomness benchmark. Nonetheless, it remains interesting to shed light on the magnitude of segregation. For example, in 2019, the Duncan index for the dimension (ii) social/occupational categories with minority group made of managers or assimilated and majority group made of the other individuals active in the labor market is estimated at 0.4. Heuristically, it means that 40% of the managers or assimilated would need to move out of their neighborhoods (defined as the LFS clusters of around thirty adjacent housing) to move in other neighborhoods in order to achieve uniform distribution of the managers or assimilated across neighborhoods among the population of active adults.

Figure 2.7 illustrates the previous discussion about the distinction of compositioninvariant and composition-variant segregation indices. Figures 2.5 and 2.6 show the estimates for two composition-invariant indices (Duncan and Theil) while Figure 2.7 presents the estimates for the Coworker index.

At first glance, the estimation suggests different conclusions regarding the magnitude and temporal evolution of segregation. However, this is only because the Coworker index captures a distinct feature of the distribution of a minority population among a majority one across units. In particular, it is sensitive to the size of the minority group. For instance, with the Coworker, the ethnicity dimension appears to have the highest magnitude of segregation, above the nationality dimension. This does not contradict the ranking given by Figures 2.5 and 2.6. It is so because the Coworker measures exposure rather than segregation, and it happens that the size of the minority group relative to the majority group is larger for ethnicity than for nationality. For example, consider non-European



Figure 2.7: same as Figures 2.5 and 2.6 for the Coworker (composition-variant) index.

ethnicity/nationality as the minority group and European ethnicity/nationality as the majority group. The proportion of the minority group within the whole population is 4% for nationalities but 16% for proxied ethnicity in 2019. Consequently, although the non-European nationalities are more concentrated in some units compared to the concentration of individual with non-European proxied ethnicity (see Figures 2.5 and 2.6), the Coworker index happens to be higher for the ethnicity dimension because the minority individuals are more numerous compared to the nationality dimension. With that specification, the Coworker is estimated around 0.22 for dimension (iv) ethnicity. We can interpret this number in the following way: if we pick an arbitrary individual with a non-European proxied ethnicity, then, in 22% of the cases, he or she will live in a neighborhood with at least one other individual with non-European ethnicity. It is around 0.17 for dimension (vi) nationality: if we pick an arbitrary individual whose nationality is non-European, then the probability that he or she lives in a neighborhood with at least one other individual with a non-European ethnicity. It is around 0.17 for dimension (vi) nationality: if we pick an arbitrary individual whose nationality is non-European, then the probability that he or she lives in a neighborhood with at least one other individual to non-European nationality is equal to 17%.

The results shown by the Coworker index about the temporal evolution and the comparison across dimensions are thus mixing the segregation by itself, understood as non-uniform distribution of minority individuals across units, and the overall size of the minority group among the majority population. That notion of exposure can also be interesting but should not be confused with the information conveyed by the composition-invariant indices.

As an illustration, the fact that the overall minority proportions are larger in dimensions (ii) and (iii) compared to dimensions (iv), (v) and (vi) explain the differences between Figure 2.7 on the one hand and Figures 2.5 and 2.6 on the other. As another final example, the increase of segregation in the sense of exposure for dimension (i) labor market positions revealed by the Coworker mainly comes from the rise of the unemployment rate during the period: the proportion of the minority group within the population has become larger.

2.4.2 Conditional analysis

In addition to the previous unconditional analysis, it is possible to study conditional segregation indices. We refer to D'Haultfœuille et al. (2021), Section 2.6 for an introduction. The LFS datasets contain various information relative to the units/neighborhoods themselves. In particular, for each unit, we know whether it belongs or not to an urban area of more than 200,00 inhabitants.

We use this information to perform a conditional analysis with, as a unit-level covariates Z, the indicator that the unit belongs to an urban area of 200,000 inhabitants. Here, such a conditional analysis amounts to performing the previous analysis separately for

- units in urban areas of more than 200,000 inhabitants (big cities), and
- units in urban areas of less than 200,000 inhabitants (smaller cities, towns, and rural areas).

After that, we can consider the segregation index restricted to one of those two subsamples of units; it can be of interest by itself. We can also compute a so-called aggregate conditional index, defined as a convex combination of the two restricted indices (whose weights are respectively the proportion of units in urban areas of more than 200,000 inhabitants among all units and the complementary proportion of units in urban areas of less than 200,000 inhabitants).

Figure 2.8 show those three types of indices plus the unconditional index:

- < 200,000 inhabitants: segregation index restricted to units in urban areas of less than 200,000 inhabitants;
- > 200,000 inhabitants: segregation index restricted to units in urban areas of more than 200,000 inhabitants;
- aggregated cond. index: convex combination of the two previous indices;
- unconditional index: segregation index estimated on all the units (as in the unconditional analysis of Section 2.4.1).

Figure 2.8: Comparison between unconditional index, aggregate conditional index, and index restricted to either units in or out of urban areas of more than 200,000 inhabitants.



There is not a unique simple conclusion of that conditional analysis. Overall, across dimensions, the estimations of the unconditional index and of the aggregate conditional

index are often close to each other (Figure 2.8 gives an illustration for dimension (i) labor market positions; Appendix 2.A presents additional figures for other dimensions). This result suggests our two main findings from the unconditional analysis (stability of segregation over time within dimension and a marked ranking of the magnitude of segregation across dimensions) remain valid in the conditional analysis.

On the other hand, in some dimensions of segregation, the estimations conducted separately on the sample of units that are part of large urban areas and on the other units reveal differences in terms of magnitude and evolution of segregation (for instance, see Figure 2.18 in Appendix 2.A for dimension (iii) education levels). Such differentiated evolution is also of interest on its own, in addition to the aggregate conditional index, to study residential segregation in-depth.

References

- Aeberhardt, R., Fougère, D., Pouget, J. and Rathelot, R. (2010), 'Wages and employment of french workers with african origin', *Journal of population economics* 23(3), 881–905.
- Allen, R., Burgess, S., Davidson, R. and Windmeijer, F. (2015), 'More reliable inference for the dissimilarity index of segregation', *The Econometrics Journal* 18(1), 40–66.
- Åslund, O. and Skans, O. N. (2009), 'How to measure segregation conditional on the distribution of covariates', *Journal of Population Economics* **22**(4), 971–981.
- Blau, P. M. (1977), Inequality and heterogeneity: A primitive theory of social structure, Vol. 7, Free Press New York.
- Carrington, W. J. and Troske, K. R. (1997), 'On measuring segregation in samples with small units', Journal of Business & Economic Statistics 15(4), 402–409.
- Cortese, C. F., Falk, R. F. and Cohen, J. K. (1976), 'Further considerations on the methodological analysis of segregation indices', *American Sociological Review* pp. 630– 637.
- D'Haultfœuille, X. and Rathelot, R. (2017), 'Measuring segregation on small units: A partial identification analysis', *Quantitative Economics* 8(1), 39–73. URL: http://dx.doi.org/10.3982/QE501
- Duncan, O. D. and Duncan, B. (1955), 'A methodological analysis of segregation indexes', American Sociological Review 20(2), 210–217.
- D'Haultfœuille, X., Girard, L. and Rathelot, R. (2021), 'segregsmall: A command to estimate segregation in the presence of small units', *The Stata Journal* **21**(1), 152–179.
- Frankel, D. M. and Volij, O. (2011), 'Measuring school segregation', Journal of Economic Theory 146(1), 1–38.
- James, D. R. and Taeuber, K. E. (1985), 'Measures of segregation', Sociological Methodology 15, 1–32.
- Massey, D. S. and Denton, N. A. (1988), 'The dimensions of residential segregation', Social forces 67(2), 281–315.
- Maurin, É. (2004), Le ghetto français: enquête sur le séparatisme social, Seuil Paris.
- Meurs, D., Pailhé, A. and Simon, P. (2006), 'The persistence of intergenerational inequalities linked to immigration: Labour market outcomes for immigrants and their descendants in france', *Population* **61**(5), 645–682.
- Rathelot, R. (2012), 'Measuring segregation when units are small: A parametric approach', Journal of Business & Economic Statistics 30(4), 546–553.
 URL: http://www.tandfonline.com/doi/abs/10.1080/07350015.2012.707586
- Reardon, S. F. and Firebaugh, G. (2002), 'Measures of multigroup segregation', Sociological Methodology 32(1), 33–67.

- Shon, J.-L. P. K. and Verdugo, G. (2014), 'Ségrégation et incorporation des immigrés en france', *Revue française de sociologie* 55(2), 245–283.
- Winship, C. (1977), 'A revaluation of indexes of residential segregation', *Social Forces* **55**(4), 1058–1066.

Appendix 2.A Additional figures

If necessary, we give below the correspondences between the French legends displayed in some of the figures and the definitions of Section 2.2.2 for each dimension and each sub-specification of minority and majority groups.

- d =acteu: labor market status dimension
 - $-m = \operatorname{actcho} (\operatorname{actif} \operatorname{chômeur}) / M = \operatorname{actocc} (\operatorname{actif} \operatorname{occupé})$: specification (a) restricted to individuals active in the labor market,
 - $-m = \operatorname{actcho} (\operatorname{actif} \operatorname{chômeur}) / M = \operatorname{actocc_inact} (\operatorname{actif} \operatorname{occupé} \operatorname{or} \operatorname{inactif}):$ specification (b) with majority group gathering inactive and active employed individuals,
- -d = cadre: social/occupational category dimension (cadre is the French word for, more or less, managers or assimilated),
 - $-m = cadre_act / M = autr_act$: specification (a) restricted to individuals active in the labor market,
 - $-m = cadre_act_ou_ancien / M = autre: specification (b) including former inactive "cadre" in the minority group,$
- d = dipl: education levels (diploma) dimension,
 - -m = bac3plus / M = autre: single specification with minority individuals having at least a bachelor's degree ("baccalauréat", French examination at the end of high shool, + three years or more) and majority individuals the others ("autre" in French),
- $d = ethni_pai_mer$: ethnicity dimension with ethnicity proxied by the countries parents' birthplaces ("pai" for "pays") with mother ("mère") as the dominant parent
 - -m = 6 / M = 1: specification (a) with Maghreb versus France proxied ethnicity,
 - -m = 67 / M = 1: specification (b) with Africa (pooling modality 6 and 7) versus France (modality 1) proxied ethnicity,
 - m = 678910 / M = tous: specification (c) with non-Europe versus Europe proxied ethnicity,
- d =**nat**: nationality dimension,
 - -m = amt / M = fr: specification (a) with "amt" standing for Algerian, Moroccan, or Tunisian and "fr" for French (same meaning for nationalities instead of birth countries of the parents that proxied ethnicity),
 - -m = afr / M = fr: specification (b) with African and French nationalities,
 - m = noneur / M = freur: specification (c), non-European and European nationalities,
 - -m = nonfr / M = fr: specification (d), non-French and French nationalities,
- d =immi: origins (immigrants) dimension,
 - $-m = \text{immi} / M = \text{autre: specification (a) with properly defined immigrants (a resident born as a foreigner abroad) and others,$
 - $-m = \text{imminatetr} / M = \text{autre: specification (b) with a subset of immigrants as the minority group, those with foreign ("étrangère") nationality.$











Figure 2.11:

Figure 2.12:





Figure 2.13:

Figure 2.14:





Figure 2.15:

Figure 2.16:







Method beta reports point-estimates Method np reports both lower and upper estimated bounds (they often coincide, resulting then in a point-estimate) Sample: units with at least 2 distinct housing, individuals with age at least 18 (23 for dimension = dipl)

Figure 2.18:





Figure 2.19:



Figure 2.20:



Index: Coworker - Units: all (smpl. if K > 50), weights: individual, method: np, (K, rho): no assumption

< 200,000 inhabitants

- > 200,000 inhabitants
- aggregated cond. index
- unconditional index





Figure 2.22:



Conditional analysis with Z > or < 200,000 inhabitants in the urban area Index: Duncan – Units: all (smpl. if K > 50), weights: individual, method: np, (K, rho): no assumption Dimension: ethni_pai_mer / 678910 / tous

Figure 2.23:



Figure 2.24:





Analysis:

< 200,000 inhabitants</p>

- > 200,000 inhabitants
- aggregated cond. index
- unconditional index
Chapter 3

Identification and estimation of a polarization index in large choice sets, with an application to U.S. Congress speech (1873-2016)

This chapter stems from an ongoing project, joint work with Xavier D'Haultfœuille and Roland Rathelot. In the continuity of Chapters 1 and 2, it deals with the identification and estimation of polarization. The difference is that it focuses on large choice sets in the sense that it is possible to observe some or even a large proportion of options that are chosen only once. In segregation, this would correspond to units with a single individual, which were rare and excluded from the analysis (remember the option withsingle).

Abstract Recently, political divisiveness appears to have been increasing in various democracies. Language, as a basic determinant of group identity, might be part of that story. "Witch hunt" versus "impeachment hearing"; "undocumented workers" compared to "illegal aliens"; "death tax" or "progressive wealth tax": those partian expressions name the same object but with different connotations. They diffuse into media coverage and can induce framing effects on public opinion. Hence the interest in measuring speech polarization of political leaders and comparing its evolution over time or across countries. A way to do so would be linguistics and literary exegesis. Another is a statistical analysis which, despite rougher (data is word counts essentially), enables the comprehensive study of a large corpus of texts without relying on ex-ante partisan expressions, and provides a measure to quantify to which extent distinct groups (e.g., Democrats and Republicans in the U.S. Congress between 1873 and 2016) speak using different words. Gentzkow, Shapiro, and Taddy (Econometrica, 2019) address this issue with a (huge) discrete choice model approach and a machine-learning type penalization. We provide an alternative method whose pros are the following: (i) a formal partial identification result for the parameter of interest (speech partisanship index) within a testable statistical model; (ii) simple and computationally light estimators for the bounds and confidence intervals; (iii) only "aggregated data" is required. As a consequence, our methodology can easily be applied to other settings with the same problem of quantifying differences as regards the choices made by individuals split into two groups in a "high-dimensional" or "large choice sets" context – meaning that the number of distinct options is large relative to the number of observed choices in data. In our application, these are choices of words pronounced by Republican and Democrat speakers, but it might as well be choices of residential locations between natives and immigrants when investigating segregation, product choices between distinct groups of consumers in empirical industrial organization, etc.

3.1 Introduction

Our goal is to reliably quantify the differences in the choices made by two exogenous groups. Various examples embody that issue. As a case in point, segregation indices assess to what extent two populations, say, natives and immigrants, make similar residential decisions. A firm might be interested to know whether two types of customers, for instance, premium subscribers versus free tier users, tend to buy different products and, if so, measure the difference. The application of this article examines speech polarization: do Republican and Democrat congresspeople speak using the same words?

A single formalization encompasses those settings. Individuals split into two exclusive groups make *choices* among a set of *options*. We wonder whether the two groups choose distinct options and how to quantify the differences between both groups as regards their choices. In our application, an option consists of a word, and the speakers choose the words they use. School segregation studies provide another illustration where an option is a class or a school, depending on the unit-level of analysis.

Qualitatively, polarization means that the options chosen by one group differ from the ones selected by the other. Several names convey that idea according to settings, such as "partisanship" in our expository application. Reasoning in terms of information further enlightens the notion. Absent polarization, the knowledge of the options chosen by an individual is uninformative about which group he or she belongs to.

We focus on situations with two groups, where each individual in the population of interest is a member of one and only one group. Moreover, groups are exogenous in the sense that membership is given ex-ante; it is not a function of the choices made. The converse problem of deducing groups from observed choices constitutes a different question that this paper leaves aside. Likewise, we do not investigate extensions to three or more groups here. Defining a concept of polarization requires extra work in such settings. To do so, convex combinations of two-group measures across all pairs of groups happen to be effective (see the literature review below). In that respect, our method could be extended to multigroup cases.

A crucial feature of many analyses devoted to group differences is that the number of possible options is large relative to the number of observed choices. Our statistical model and asymptotics specify this idea of high-dimensional choice sets.

Such a context logically arises whenever the options display some capacity constraints, in so far as the number of choices per option is capped. Classes with limited numbers of pupils epitomize that framework: if the number of students grows to infinity, so does the number of classes. Besides, the relevance of asymptotics in the number of options can expand into environments where there are no intrinsic capacity limitations. A significant example concerns speech or text data, as in our leading application. Indeed, Herdan's law (Remark 3.1) reveals that, empirically, the vocabulary size of any corpus of texts increases with its length. Expressed in our theoretical terminology, the number of options is an increasing function of the number of choices.

Furthermore, the dimension, or size, of any choice set, namely the number of options, depends upon the exact definition of options. Researchers may have some latitude in the specification. This leeway gives another argument to underpin our asymptotics. Provided adequate data quality, researchers naturally would want to consider narrower, more precisely delimited options as they get additional observations, that is, as they record more choices. To take advantage of a famous example, imagine a study of travel modes contrasting men's and women's decisions. A few data points drive to restrict to cars versus

buses, whereas larger datasets allow dividing the latter alternative into blue buses and red buses, hence increasing the number of options.

We propose reliable measures in several senses. First and foremost, previous literature showed that high-dimensional choice sets cause a "small-unit bias" of naive methods; those that quantify group differences through the dispersion, across options, of the shares of choices made by one of the groups. That issue hampers meaningful comparisons over time or, for instance, across countries since the magnitude of the bias might change too. Our method addresses and corrects the bias. Therefore, it suits the common framework of a large number of options and enables trustworthy comparisons of polarization measures.

Secondly, formal identification, estimation, and inference results support our method. We derive them in a statistical model with two appealing characteristics. One of its central assumptions is testable, meaning that it is possible to reject it if the implied data-generating process cannot rationalize the observations. Also, the asymptotics in options growing to infinity fits with high-dimensional choice sets. The parameter of interest is a polarization index. In our framework, we prove that, in general, it is partially identified due to the presence of options that receive a single choice. Specifically, the length of the identification interval is proportional to the fraction of options chosen only once. The intuitive explanation is the following: for a given option selected only once, we are unable to determine whether this option, as compared to the other options, is more or less chosen by one group instead of the other. Our identification theorem thus shapes a fundamental trade-off between, on the one hand, how much information on the intensity of polarization can be retrieved from the data and, on the other, how precisely defined can be the options. We provide consistent estimators for the bounds of the identification interval, as well as confidence intervals for the index with guaranteed asymptotic coverage. Additionally, we suggest an extrapolation strategy that, under a supplemental regularity condition, yields a point-estimate of the polarization index.

A final aspect of the reliability of our approach deals with its practical implementation. Regarding computational requirements, our method does not entail any numerical approximation and is efficiently executed in standard machines. Indeed, empirical means alone form our estimators and confidence intervals. As a consequence, their computation is straightforward and virtually immediate, even in huge databases.¹ Last but not least, our method is economical concerning the type of data needed. To make use of a familiar Industrial Organization parallel, like Berry, Levinsohn and Pakes (1995), we only require "aggregated" data, namely, the number of choices for each option by groups, in contrast to individual choices. In other words, keeping track of the identity of choosers is unnecessary. Still, in some situations, it is possible to do so, and individual characteristics beyond group-membership may be available. This is the case when studying speech polarization in the U.S. Congress in particular. The interest in such information stems from the fact that an individual's choices are not entirely determined by his or her group but conjointly by personal attributes. Hence, in complement to a raw, or unconditional, index of polarization, researchers might want to measure to what extent two groups make distinct choices net of the influence of additional characteristics. Let conceive a simplified illustration in our context. Democrat and Republican speakers could use different words not only because of their opposite party affiliations but owing to men and women having distinct languages and women being more numerous among Democrats than Republicans. We propose an extension of our method to obtain conditional indices.

¹For instance, a laptop estimates the identification bounds in a database of over two million options, as it is the case in our application, in a couple of seconds.

The paper is organized as follows. Section 3.2 defines our statistical model and the parameter of interest. In Section 3.3, we present our main results: the sharp identification of the bounds on our index and a related confidence interval as well as an extrapolated point-estimator. Section 3.4 describes two extensions, namely the inclusion of covariates to define conditional polarization indices and the test of the binomial assumption that underlies our method. Section 3.5 applies our approach to investigate speech polarization between Republican and Democrats congresspeople from 1873 to 2016. In doing so, we revisit Gentzkow, Shapiro and Taddy (2019), who considered initially that question and developed a new method to measure "group differences in high-dimensional choices." Section 3.6 concludes and discusses possible limits of our approach. A literature review completes the introduction and discusses distinctions and similarities between their method and ours. Appendix 3.A displays the proofs of the main theorems. Appendix 3.B presents supplemental material regarding our methodology. Appendix 3.C provides further details on the application. Appendix 3.D establishes a formal link between our model and that of Gentzkow et al. (2019).

Literature In a broad perspective, note a divergence between the class of discrete choice models, triggered by McFadden (1974), and the settings we consider: the former holds fixed the number of options while the point of high-dimensional choice sets is the asymptotics in both the numbers of options and observed choices.

In such contexts, segregation indices are prominent attempts to quantify differences between the choices made by two groups. Despite lack of formal asymptotics or sampling design, Cortese et al. (1976) and Winship (1977) unveil the existence of a "small-unit" bias of historical constructions of indices, such as the dissimilarity index of Duncan and Duncan (1955) or classical indices surveyed in James and Taeuber (1985). Here, the term "unit" refers to what we name options. The epithet "finite-sample" is sometimes used to denote the bias. We do not endorse the expression because it insinuates larger datasets are the solution. Nevertheless, the problem is deeper as the numbers of options and choices grow simultaneously. The core of small-unit bias is that, as we observe few choices per option, it is frequent in the data to see options that are, by chance, chosen by only one of the two groups. This may wrongly suggest those options are elected by one group but never the other, thus pointing towards high polarization, whereas the underlying probabilities to choose those options can be the same for both groups. Carrington and Troske (1997) and Allen et al. (2015) propose techniques to estimate the magnitude of the bias in order to correct it. Our approach is closer to Rathelot (2012) and D'Haultfœuille and Rathelot (2017) in so far as, instead of an ex-post correction, we define the index as a function of the choice probabilities, which are arguably the relevant primitives of interest to measure polarization.² Although the article does not describe the extension of our method to situations with more than two groups, we mention Reardon and Firebaugh (2002) as a seminal reference for multigroup segregation indices.

Alongside the literature on segregation, our primary inspiration is Gentzkow, Shapiro and Taddy (2019) (GST). They confront the very same issue of quantifying differences between the decisions made by two groups in large choice sets. They develop a pioneering method to do so and apply it to study partian differences in language in the U.S. Congress from 1873 to 2016. Their results highlight that accounting for the bias is essential to uncover the evolution of partianship over time. GST shows that after a century-long

²We refer to Winship (1977) for further discussion of segregation indices and the opposition between "evenness" and "randomness" benchmarks to appraise group differences.

period of low and stable speech polarization, it surged from the 1990s and seemed to plateau in recent years. Thorough analyses explore credible explanations and the features of the increase, notably distinguishing within- versus between-topic partisanship. Nonetheless, we think our article brings additional contributions.

The bulk concerns methodology since we believe our method is an alternative to GST's in settings of high-dimensional choice sets.³ To sum up, GST uses a generative model of speech linked to a discrete choice model. The latter expresses the utility for a speaker with given personal characteristics and party to choose each word of a vocabulary that forms the choice set. GST constructs a penalized estimator to infer the parameters of the model, from which they define a partisanship index. The Lasso-type, or L_1 , penalization on the coefficient capturing the effect of party affiliation is critical to handle the bias.

To weigh both methods, we draw upon an analogy of the two estimators involved in Hausman's tests (Hausman, 1978). We would say our method is more robust in the sense that it requires aggregated data and yields tractable estimators, which are backed by identification, estimation, and inference results obtained with asymptotics in the number of options. In contrast, GST's would be the precise but perhaps more specific estimator as, in our understanding, it necessitates several elements, possibly demanding in the settings we consider. First, it uses thinner observations, namely choices by individuals, instead of totals by groups. Second, the theoretical guarantees of inference, performed via subsampling, are derived under fixed vocabulary and the number of unique speakers tending to infinity. The conditions appear at odds with the idea of high-dimensional choice sets.⁴ Besides, those guarantees pertain to an estimator that GST's method only approximates. A large number of options indeed complicates the numerical estimation of their multinomial logit model, and, in their application, they resort to a Poisson approximation. Incidentally, the procedure entails setting some hyper-parameters, although GST explains how to determine the penalty loads. On the other hand, GST is at an advantage to incorporate personal characteristics of choosers, even chooser-level random effects, as in their Supplemental Material. They easily include continuous covariates, whereas it is hard to do so in our approach, as our definition of conditional indices (Section 3.4.1) reveals. Hence, GST's method is likely better at extracting the specific influence of group-membership among diverse individual factors. The discrete choice model defined at individual-level allows that flexibility but costs GST suitability for high-dimensional choice sets: asymptotics in the number of individuals while options are fixed; complex and numerically challenging estimation to deal with a large number of alternatives. In comparison, we tailor our method to such settings where the elementary objects are options and groups. Consequently, it is less adapted to individual-level analyses and, when choosers' identities are known, arguably rougher, but in return, steadier. All in all, we are confident the two approaches are complementary, depending on the context and available data.

Another distinction bears on the definition of polarization index in relation to the group proportions within the population. As far as we understand, the partisanship index of GST takes as a benchmark to quantify polarization, a situation where the two

³GST presents three estimators: plug-in or maximum likelihood, leave-out, penalized. We discuss the third one, defined in §4.3, p1316, and labeled by GST their "preferred" estimator.

⁴GST writes, p1318, "Though we do not pursue formal results for the case where the vocabulary grows with the sample size, we note that such asymptotics might better approximate the finite-sample behavior of our estimators." Our approach builds on this idea. Furthermore, using that asymptotics of a choice set (vocabulary) growing with observed choices (sample size), Appendix 3.D formally brings GST's framework and ours together.

groups halve the population. They consider an observer trying to deduce an individual's group from one of his or her choices. The observer starts from a "neutral prior" (GST, p1313), assigning an equal probability of 1/2 to both groups. However, in most cases, there is no reason that the two groups share the same proportion. We think there exist genuine interrogations on what we want to pay heed to when contrasting decisions made by two groups. Even when the proportions are not half-half, a neutral prior might be sensible in sharp bipartite oppositions, like Democrats versus Republicans in U.S. politics. Nonetheless, we advocate that, in general, polarization should refer to a concept of dispersion alone without consideration for differences in the number of choices made by the two groups. In our reading of the segregation indices literature, the desired property of composition invariance supports that viewpoint (see, for instance, James and Taeuber (1985), p15). That is why we define a polarization index that focuses on dispersion and does not posit any prior composition of the population. In that new respect, our approach could be deemed more robust than GST's.

The second line of contributions concerns the application to speech partisanship. For any research processing text data to study polarization, our analysis underscores that the specification of the dictionary matters. Section 3.5.2 provides details on the processing operations that lead to the set of words included in the sample.⁵ We adhere to GST's steps except for selections based on the number of occurrences. GST drops phrases that are not sufficiently often pronounced in the data. Within our framework, our angle mindful to identification reveals it is those restrictions that buy most of the identification power.⁶ Interestingly, absent this form of selection, two-year-long Congress session transcripts exhibit a substantial fraction of phrases pronounced only once, over 60% for the couples of words as processed by GST. As a consequence, without such restriction, our identification interval is barely informative and cannot reject a constant polarization throughout the 1873-2016 period. Some processing operations and selections may, of course, be sensible and necessary. However, we believe our approach is helpful to gauge the impact of the restrictions, especially in terms of identification power.

With selection on the number of occurrences or using our extrapolated point-estimators, our broad conclusion is akin to GST's findings. We identically detect a dramatic increase from the 1990s. A closer look shows our estimates come slightly at variance with GST's analysis, for they temper the picture of a flat and low level of polarization until that unprecedented contemporary climb. Though we measure the highest polarization in recent years, we observe periods of comparable intensity in past U.S. history, notably in the first decade of the twentieth century and, to a lesser extent, during the twenties and fifties.

3.2 The framework

We state the set-up in the context of speech partial performance between Democrat and Republican congresspeople. As the Introduction emphasized, our approach is more general, and we sometimes widen the presentation and refer to a general setting.

⁵In fact, an option consists of a *bigram*, a couple of words, also denominated a *phrase*, since we follow GST with a popular N-gram model of text, taking N = 2. Our methodological point remains.

⁶Selecting words above some thresholds of occurrences mechanically decreases the proportion of words included in the analysis that are pronounced only once, hence narrowing our identification interval.

3.2.1 Statistical model

In a given period, for instance, a two-year Congress session, we observe the speeches pronounced by speakers, who are split into two groups, Democrats and Republicans. In our particular application, texts, with greater reason, speeches, are complex data. Various perspectives could be relevant to tackle the issue. Here, a "bag-of-word" modeling (see Section 3.5 below) enables to transform the raw transcripts into counts of words. The operation shapes initial text data into a generic form in which words, that is, options in our abstract terminology, constitute the basic statistical units.

For each word (option) indexed by j among a vocabulary (choice set) $\mathcal{J} := \{1, \ldots, J\}$, $J \in \mathbb{N}^*$, we observe K_j^R the number of occurrences (choices) pronounced by Republicans (made by members of a reference group, arbitrarily chosen out of the two); respectively K_j^D for Democrats. We define $K_j := K_j^R + K_j^D$ the total number of occurrences of the word.

Hence, for each word j, K_j^R/K_j is the observed proportion of occurrences of the word said by Republicans. A natural idea to assess to what extent the two groups choose distinct words is to look at the variation in the proportions $\{K_j^R/K_j\}_{j\in\mathcal{J}}$ across words. If they are all roughly equal, Republicans and Democrats more or less use the same words. On the contrary, in an extreme case where those proportions are equal to either 0 or 1, the two groups speak entirely distinct languages.

Although sensible, that approach suffers from a small-unit bias when the numbers $\{K_j\}_{j\in\mathcal{J}}$ of choices per option are small; that is when the choice set is large relative to the number of observed choices. To address the problem, we follow Rathelot (2012) and D'Haultfœuille and Rathelot (2017) and introduce an underlying unobserved object. For each word j, ρ_j is the probability, conditional on word j being pronounced, that it is said by a Republican. The share K_j^R/K_j is a straightforward estimator of ρ_j . However, for finite, a fortiori small, $\{K_j\}_{j\in\mathcal{J}}$, the variation in $\{K_j^R/K_j\}_{j\in\mathcal{J}}$ could merely be due to random variability while the $\{\rho_j\}_{j\in\mathcal{J}}$ are all equal. That is why naive indices based on observed proportions overestimate the real level of polarization, which refers to variations in the probabilities $\{\rho_j\}_{j\in\mathcal{J}}$ across options.

Our assumption about the data-generating process (DGP) is the following.⁷

Assumption 3.1 (DGP). We observe an i.i.d. sample $(K_j^R, K_j)_{j=1,...,J}$ such that (K_j^R, K_j, ρ_j) has the same distribution as (K^R, K, ρ) , which satisfies: $\mathbb{E}[K] > \mathbb{E}[K\rho] > 0$, and

$$K^R | K, \rho \sim \text{Binomial}(K, \rho).$$
 (3.1)

This sets a binomial mixture model. The model is non-parametric, the distribution of K and ρ being unspecified, and involves a latent variable, ρ . In our application, the i.i.d. hypothesis formalizes the implications of N-gram model, which converts texts into counts of words considered as unrelated. In other settings, sampling-based modeling that assumes the data consists of a random sample of options will usually rationalize the assumption. The technical condition $\mathbb{E}[K] > \mathbb{E}[K\rho] > 0$ only says that we do observe some choices from both groups.

The other part of our DGP stipulates the distribution of K^R conditional on K and ρ . In a sense, it expands the independence inter-options stemming from i.i.d. samples to the independence of choices intra-options. Indeed, it precludes interactions across the choices

⁷As for now, we focus on our unconditional analysis; Section 3.4.1 introduces conditional indices.

falling in the same option since repetitions of Bernoulli trials form a Binomial distribution. In congressional speeches, for instance, it posits that the fact that a particular word is pronounced at the beginning of the session by, say, a Democrat, does not influence the party affiliation of the speaker who will pronounce the next occurrence of the word. How restrictive the condition is depends on the context. In any case, the model permits to test it: Section 3.4.2 presents a test whose null hypothesis is the binomial assumption. The test is consistent and has asymptotic level guarantees.

As the Introduction explained, we consider the asymptotics where J tends to $+\infty$ in order to suit high-dimensional choice sets. In our application, two arguments support that reasoning. First, asymptotics can be interpreted as an approximation, and the vocabulary contains millions of phrases in our case. Second, it complies with the way we talk and write in practice, as revealed by Herdan's law.

Remark 3.1 (Herdan's, or Heaps's, law). *Herdan (1960) in a linguistics perspective, or Heaps (1978) in statistical text analysis, highlight an empirical relationship between the number J of words ("types" in a text analysis terminology) and the number n of occurrences ("tokens") in a corpus of texts:*

$$J = \beta n^{\gamma},$$

for two constants, $\beta > 0$ and $\gamma \in (0,1)$, that depend on the genre. It happens that numerous English corpora, such as Shakespeare's plays, Brown corpus, Google N-grams, satisfy the relation and display γ ranging from 0.35 to 0.75. Thus, the vocabulary size J of a corpus of texts typically increases with its length n. Figure 3.13 in Section 3.C.2 evidences that fact for congressional transcripts.

3.2.2 Parameter of interest

Polarization measures, like segregation indices in James and Taeuber (1985), were historically defined as inequality indices on the observed proportions $\{K_j^R/K_j\}_{j\in\mathcal{J}}$, hence blurring the distinction between population estimands and sample quantities. In contrast, we follow a common modern approach. We define parameters of interest as population quantities, formally features of the distribution of the data, for we place ourselves in an infinite population modeling. Then, we estimate them with sample counterparts, using the analogy principle (Manski, 1988).

The second consideration leading to our estimand relates to the small-unit bias or, more appropriately, the proper notion of differences in choices between groups. Guided by Winship (1977), we adopt a "randomness" benchmark to quantify them. It means polarization is not characterized by differences in the observed choices, which cannot be more than particular sample realizations, but essentially by the differences in the underlying probabilities, which entail systematic divergences regarding how both groups make choices. The former are informative on the latter, and since the probabilities remain unobserved, we cannot help but use the observed choices for estimation. However, our parameter relies on probabilities.

Those two points justify defining our estimand as a function of P^{ρ} , the distribution of ρ .⁸ On the other hand, they say nothing about the form of the function. The issue has generated long-standing debates on the properties and desirable axioms of competing indices. This paper does not enter that field because we stick to the measure of polarization

⁸For any random variable V (or tuple of random variables), P^V denotes its (joint) distribution.

studied by Gentzkow et al. (2019). GST specifies a particular mapping from P^{ρ} to define their partisanship index. We use the same mapping subject to adaptation to our setting and generalization to groups with unequal proportions.⁹

A heuristic path toward our index begins with $\pi' := 1 - 2\mathbb{E}[\rho(1-\rho)]$. If ρ follows a Bernoulli distribution, π' is maximal equal to 1: the intersection of options chosen by both groups is void, and polarization is complete. On the contrary, if ρ follows a Dirac distribution, being almost surely equal to $q \in (0, 1)$, then $\pi' = 1 - 2q(1-q)$ and is higher as q is further from 1/2. In this case, given the definition of ρ , q corresponds to the share of speech from Republicans, that is, whatever the word pronounced, the probability that a Republican says it equals q. In the Introduction, we contend that, in general, polarization should refer to a concept of dispersion only. Hence, our final estimand adjusts π' to ensure that whenever P^{ρ} is a Dirac, our polarization index is minimal, hence coherent with no difference in the probabilities between groups.

Our index differs from π' in another respect. Under Assumption 3.1, the joint distribution of (K, ρ) entirely determines the model. Logically, our estimand is a function of $P^{(K,\rho)}$. The number of choices per option K indicates the popularity of options. A priori, K and ρ are related. In our application, for instance, if speech partial priori matters and is leveraged by politicians, it is reasonable to think that the most partian words $(\rho \text{ close to } 0 \text{ or } 1)$ are also the most frequent (large K). Our framework is agnostic on possible links between K and ρ . The inclusion of K weighs options proportionally to their popularity, giving more weights to options that receive more choices. Other weights are conceivable and would affect the measure of polarization.¹⁰ Nevertheless, we think our weighting scheme is transparent and intuitive. Imagine a toy example where each group accounts for half of the choices and, for all options except one, the probabilities ρ are equal to 1/2 while the probability associated with the unique divisive option is, say, 0.03, that is, when chosen, that option is chosen by virtually one of the two groups only. A polarization index condenses in a scalar number the differences in choices between groups. In this situation, it is natural to describe the level of polarization as higher as the frequency of the last option, meaning the proportion across all choices of decisions for that option gets larger. Our weights does just that, transforming π' into $\pi'' := 1 - 2\mathbb{E}[K\rho(1-\rho)]/\mathbb{E}[K]$.

Finally, our index generalizes π'' to unequal group proportions and writes:

$$\pi := 1 - \frac{\mathbb{E}[K\rho(1-\rho)]}{2\mathbb{E}[K]\,q(1-q)}\,,\tag{3.2}$$

where $q := \mathbb{E}[K\rho]/\mathbb{E}[K]$. The binomial condition in Equation (3.1) and the law of iterated expectations imply $q = \mathbb{E}[K^R]/\mathbb{E}[K]$ Thus, q is the share of speech pronounced by Republicans. In general settings, it is the proportion of choices made by individuals from the reference group. When individuals make one choice each, it boils down to the proportion of individuals in the reference group.

The index π is well-defined by (3.2) provided that there are some choices of course, $\mathbb{E}[K] > 0$, and $q \in (0, 1)$ (or, equivalently, $\mathbb{E}[K] > \mathbb{E}[K\rho] > 0$). Otherwise, there is only

 $^{^{9}}$ Appendix 3.D makes a formal link between our estimand and GST's. GST also defines their measure of partisanship as a function of underlying probabilities, in contrast to observed choices.

¹⁰D'Haultfœuille and Rathelot (2017) (Appendix B.1) discuss unit- versus individual-level weights for segregation indices. Another possibility here could assign the same weight to each option, irrespective of their frequencies.

one group, and a notion of polarization is not applicable. The following result proves that π is a meaningful measure of polarization.¹¹

Theorem 3.1 (Polarization index). Assume $\mathbb{E}[K] > \mathbb{E}[K\rho] > 0$. Then (i) $\pi \in [1/2, 1]$; (ii) $\pi = 1/2$ if and only if $\rho | K > 0$ follows a Dirac distribution; (iii) $\pi = 1$ if and only if $\rho | K > 0$ follows a Bernoulli distribution.

3.3 Main theoretical results

We consider hereafter an asymptotic set-up in J. In this framework, the joint distribution of (K^R, K) is identified. We first derive simple bounds on π under Assumption 3.1 only and show how to conduct inference based on these bounds. We then consider an extrapolation strategy allowing us to obtain a point estimator of π .

3.3.1 Bounds under minimal conditions

The following theorem gives sharp bounds on π under Assumption 3.1 only and underscores the source of partial identification.

Theorem 3.2 (Identification). Suppose Assumption 3.1 holds. Then π is partially identified, with lower and upper bounds satisfying

$$\begin{split} \underline{\pi} &:= 1 - \frac{\mathbb{E}[K]}{2\mathbb{E}[K^R]\mathbb{E}[K^D]} \left\{ \mathbb{E}\Big[\frac{K^R K^D}{K-1} \mathbbm{1}\{K > 1\}\Big] + \frac{\mathbb{E}\Big[K^R \mathbbm{1}\{K = 1\}\Big] \mathbb{E}\Big[K^D \mathbbm{1}\{K = 1\}\Big]}{\mathbb{E}\big[\mathbbm{1}\{K = 1\}\big]} \right\},\\ \overline{\pi} &:= 1 - \frac{\mathbb{E}[K]}{2\mathbb{E}[K^R]\mathbb{E}[K^D]} \mathbb{E}\Big[\frac{K^R K^D}{K-1} \mathbbm{1}\{K > 1\}\Big], \end{split}$$

where we use the convention that $\mathbb{E}[K^R \mathbb{1}\{K=1\}]/\mathbb{E}[\mathbb{1}\{K=1\}] = 0$ when $\mathbb{P}(K=1) = 0$, and $\mathbb{1}\{K>1\}/(K-1) = 0$ if K = 1. Moreover, the bounds are sharp. Finally, $\underline{\pi}$ and $\overline{\pi}$ do not depend on $\mathbb{P}(K=0)$. Hence, they are also the sharp bounds if only the distribution of (K, K^R) conditional on K > 0 is known.

The length Δ of the identification interval satisfies

$$\Delta = \mathbb{E}\left[K^R \mid K=1\right] \mathbb{E}\left[K^D \mid K=1\right] \mathbb{P}(K=1).$$
(3.3)

Hence, the bounds become tighter as $\left|\mathbb{E}\left[K^{R} \mid K=1\right] - 1/2\right|$ increases and as $\mathbb{P}(K=1)$ decreases. Options that are chosen only once constitute the main source of partial identification. The reason is that contrary to $m(k) := \mathbb{E}(\rho^{2} \mid K=k)$ for k > 1, $m(1) := \mathbb{E}(\rho^{2} \mid K=1)$ is not point identified, but can simply be bounded using Jensen's inequality and $\rho^{2} \leq \rho$:

$$\mathbb{E}(K^R \mid K = 1)^2 = \mathbb{E}(\rho \mid K = 1)^2 \le m(1) \le \mathbb{E}(\rho \mid K = 1) = \mathbb{E}(K^R \mid K = 1).$$
(3.4)

¹¹Polarization indices, such as segregation indices, are often normalized between 0 and 1. Here, following GST, the normalization operates in [1/2, 1]. Appendix 3.B.1 shows that π , via an affine transformation, is linked to a classical residential segregation index named, among others denominations, the Coworker index.

Theorem 3.2 is related to, but distinct from, the small-unit bias issue, discussed in our context by Gentzkow et al. (2019) but also by Winship (1977) and Carrington and Troske (1997), among many others, in the measure of segregation. In our context, the small-unit issue refers to the (positive) bias of the naive estimator of π based on replacing ρ by K_R/K in Equation (3.2). This positive bias disappears as K tends to infinity. However, Theorem 3.2 reveals that it is not necessary to have K tending to infinity for π to be point identified: $\mathbb{P}(K = 1) = 0$ is enough for that purpose. But of course, the corresponding estimator is different from the aforementioned naive estimator.

The last point of the theorem states that the bounds are independent of $\mathbb{P}(K = 0)$ or, equivalently, are also the sharp bounds if one only identifies the distribution of (K, K^R) conditional on K > 0. This is fortunate because knowing $\mathbb{P}(K = 0)$ requires to take a stand on the whole set of options, including those never chosen, and this may not be obvious. In the case of text data, one could consider words from a given dictionary, but people may use foreign words or new words that are not in the dictionary yet. Similarly, one may wonder whether it makes sense to include some old words that are not used anymore.

The expression of the identification bounds in Theorem 3.2 suggests the following plug-in estimators:

$$\hat{\underline{\pi}} := 1 - \frac{\sum_{j=1}^{J} K_j}{2\sum_{j=1}^{J} K_j^R \sum_{j=1}^{J} K_j^D} \times \left[\sum_{j=1}^{J} \frac{K_j^R K_j^D}{K_j - 1} \mathbb{1}\{K_j > 1\} + \frac{\sum_{j=1}^{J} K_j^R \mathbb{1}\{K_j = 1\} \sum_{j=1}^{J} K_j^D \mathbb{1}\{K_j = 1\}}{\sum_{j=1}^{J} \mathbb{1}\{K_j = 1\}} \right],$$

$$\sum_{j=1}^{J} K_j = \frac{J}{K_j} K_j^R K_j^D$$

$$\widehat{\pi} := 1 - \frac{\sum_{j=1}^{J} K_j}{2\sum_{j=1}^{J} K_j^R \sum_{j=1}^{J} K_j^D} \sum_{j=1}^{J} \frac{K_j^R K_j^D}{K_j - 1} \mathbb{1}\{K_j > 1\},$$
(3.5)

where we use here similar conventions as those in Theorem 3.2, but with sample means instead of expectations. With such conventions, $\hat{\pi} = \hat{\pi}$ if $K_j \neq 1$ for all j.

A straightforward application of the Central Limit Theorem ensures that $(\hat{\pi}, \hat{\pi})$ are asymptotically normal. Similarly, their asymptotic variances $\underline{\omega}$ and $\overline{\omega}$ can be estimated consistently by plug-in estimators, which we denote by $\hat{\underline{\omega}}$ and $\hat{\overline{\omega}}$. We refer to Propositions 3.1 and 3.2 in Appendix 3.B.3 for further details and focus here on the construction of confidence intervals (CI) for π .

To do so, we follow Imbens and Manski (2004) and Stoye (2009). For any $\alpha \in (0, 1)$, we define our CI by

$$CI_{1-\alpha}^{\pi} := \left[\underline{\widehat{\pi}} - \frac{q(\alpha)\sqrt{\widehat{\omega}}}{\sqrt{J}}, \overline{\widehat{\pi}} + \frac{q(\alpha)\sqrt{\widehat{\omega}}}{\sqrt{J}} \right], \qquad (3.6)$$

where $q(\alpha)$ solves

$$\Phi\left(\mathbf{q}(\alpha) + \frac{\sqrt{J}\widehat{\Delta}}{\max\left\{\sqrt{\widehat{\omega}}, \sqrt{\widehat{\omega}}\right\}}\right) - \Phi(-\mathbf{q}(\alpha)) = 1 - \alpha, \tag{3.7}$$

with $\widehat{\Delta} := \widehat{\pi} - \widehat{\pi}$ and $\Phi(\cdot)$ the cumulative distribution function of the standard normal distribution. Under technical conditions (see Proposition 3.3 in Appendix 3.B.3), this confidence interval is asymptotically valid uniformly over a set of DGP including point-identified cases.

3.3.2 Point-identification and estimation through extrapolation

In practice, the length Δ of the identification interval can be large. Equation (3.3) shows that this is the case if $\mathbb{P}(K = 1)$ is large and $\mathbb{E}[K^R | K = 1] \simeq 1/2$. These conditions are not unlikely: in our application below, $\mathbb{P}(K = 1) \simeq 60\%$ and the second condition also holds. We now investigate a way to achieve point identification based on an extrapolation argument. First, remark that

$$\pi = 1 - \frac{\mathbb{E}[K]}{2\mathbb{E}[K^R] \mathbb{E}[K^D]} \left\{ \mathbb{E}\left[K^R\right] - \mathbb{E}\left[\frac{K^R(K^R - 1)}{K - 1}\mathbb{1}\{K > 1\}\right] - \mathrm{m}(1)\mathbb{P}(K = 1) \right\}, \quad (3.8)$$

where we recall that $m(k) = \mathbb{E}[\rho^2 | K = k)$. This equation highlights the sole source of partial identification of π , namely m(1). Then, our bounds on π are based on the inequalities (3.4) on m(1). These inequalities do not exploit any information on $(m(k))_{k\geq 2}$. Yet, in practice, $k \mapsto m(k)$ is likely to be regular. For instance, $k \mapsto m(k)$ is constant if $\rho \perp K$. More generally, we suggest the following strategy:

- 1. assume a parametric model on $(m(k))_{k \in \{1,...,\overline{k}\}}$ for some $\overline{k} \ge 2$: $m(k) = f(k, \theta_0)$ for a known function f and an unknown θ_0 ;
- 2. identify θ_0 using $(m(k))_{k \in \{2,...,\overline{k}\}}$;
- 3. recover m(1) as $f(1, \theta_0)$.

The second step requires that the map $\theta \mapsto (f(2,\theta), ..., f(\overline{k},\theta))$ be injective. Hence, we face a trade-off here between the flexibility of the model (i.e., the dimension of θ is large) and the use of options that are close enough to options chosen only once (i.e., \overline{k} remains small). Examples of parametric models are polynomial functions, with $f(k,\theta_0) = \sum_{j=0}^d \theta_{0j} k^j$ with $\theta_0 = (\theta_{00}, ..., \theta_{0d})$. In this case, θ_0 is identified if and only if $\overline{k} \ge d + 2$. This extrapolation idea is close to that used in regression discontinuity designs (RDD) with discrete running variables, see in particular Lee and Card (2008), Kolesár and Rothe (2018). As in such designs, the parametric model should only be seen as an approximation of the true function $k \mapsto m(k)$.

To implement that strategy and construct an extrapolated point-estimator, we start by defining, for any integer k > 1 observed in the data, that is $\sum_{j=1}^{J} \mathbb{1}\{K_j = k\} \neq 0$, the following estimator for m(k)

$$\widehat{\mathbf{m}(k)} := \frac{\sum_{j=1}^{J} K_{j}^{R} (K_{j}^{R} - 1) \mathbb{1}\{K_{j} = k\}}{k(k-1) \sum_{j=1}^{J} \mathbb{1}\{K_{j} = k\}}.$$
(3.9)

Then, we require a specific functional form to extrapolate m(1). We focus on polynomial extrapolation from the support $k \in \{2, 3, \ldots, \overline{k}\}$.¹² For a polynomial order $r \in \mathbb{N}$ and a closing integer $\overline{k} \geq r + 3$, we introduce the null hypothesis $H_0(r, \overline{k})$: "on $\{2, 3, \ldots, \overline{k}\}$, $k \mapsto m(k)$ is a polynomial function of order r". Under that hypothesis, we are in the setting of Classical Minimum Distance (CMD) Estimation. With the terminology of Wooldridge (2010), §14.5, the vector $\widehat{[m(k)]}_{k=2,\ldots,\overline{k}}$ estimates the reduced-form parameters $[m(k)]_{k=2,\ldots,\overline{k}}$,

¹²Other specifications are feasible. The estimates $\{\mathbf{m}(k)\}_{k=2,3,...}$ can be used to assess a credible form for the function $k \mapsto \mathbf{m}(k)$ on some domain $\{2, 3, \ldots, \overline{k}\}$. Polynomials are probably enough in most settings. The restriction to small values for k to perform extrapolation is logical as we target $\mathbf{m}(1)$. Besides, a lack of data on that range suggests a large-unit setting, with K above some threshold. It would be weird to observe a large fraction of K = 1 in such situations, so that π is likely point-identified or, at least, the length of the identification interval is small and sufficiently informative, hence the irrelevance of an extrapolated estimator.

and, through the polynomial assumption, there is a known map from the structural parameters, namely the coefficients of the polynomial, to $[m(k)]_{k=2,...,\overline{k}}$. Two results follow. First, we have an efficient CMD estimator of the coefficients of the polynomial, and their sum gives an estimator $\widehat{m(1)}$ of m(1) by extrapolating the polynomial to k = 1. Second, an over-identification test of the restriction $H_0(r, \overline{k})$ is available, and allows to check the chosen functional form is compatible with the data.

Given m(1), an extrapolated point-estimator of π is obtained by plug-in in Equation (3.8), swapping expectations for means, $\hat{\pi}_{\text{extrapolated}} :=$

$$1 - \frac{\sum_{j=1}^{J} K_j}{2\sum_{j=1}^{J} K_j^R \sum_{j=1}^{J} K_j^D} \Biggl\{ \sum_{j=1}^{J} K_j^R - \sum_{j=1}^{J} \frac{K_j^R (K_j^R - 1)}{K_j - 1} \mathbb{1}\{K_j > 1\} - \widehat{\mathbf{m}(1)} \sum_{j=1}^{J} \mathbb{1}\{K_j = 1\} \Biggr\}.$$
(3.10)

3.4 Extensions

3.4.1 Including covariates

Group membership is likely not the unique determinant of the choices made. Individual characteristics can also matter, like age, gender, constituency's location for the language congresspeople use. In some settings, researchers observe individual choices or, at least, aggregated counts broken down by group membership plus choosers' covariates, as expressed in Assumption 3.2. Conditional polarization indices aim at quantifying the differences in the choices made between two groups net of the influence of additional factors. The objective is analogous to controls in linear models, where a specification includes covariates to monitor the variability of the outcome. In conjunction with an unconditional measure, conditional indices of polarization enable to disentangle the sources of group differences.

Orthogonal to the effects of individual features, it can be interesting to unpack polarization by different subsets of the choice set. In our application, following GST, we can wonder whether the level of speech partisanship varies with the debated topic; for example, are the differences in language between Republicans and Democrats larger when they discuss budget or immigration? To do so, one can gather the words into different types according to their topics. Here, covariates are defined at the level of options.

Below, we present an elementary extension of our framework to conditional indices, the idea being to consider convex combinations of within-type raw indices. We distinguish individual- from option-level covariates since they have different implications. The optionlevel case works fine, whereas, as discussed in the Introduction, our approach faces some difficulties for rich individual-level characteristics compared to GST's method.

Options' characteristics We restrict to discrete covariates and suppose they are synthesized into a single categorical variable X with values in $\{1, \ldots, \overline{X}\}, \overline{X} \ge 2$, which indicates the type of an option. For instance, in our application, it can be a partition of words into topics, like budget, defense, health, immigration. For each topic, Republicans and Democrats may use the same language to talk about it but differ on their distribution of speaking time across topics. On the whole, it would lead to differentiated language between the two parties, that is, speech polarization, but stemming from between-topic instead of within-topic partisanship. To unravel the distinction, a simple idea is to study polarization indices within each type and then aggregate them into a summary conditional measure.

To do so, we assume we observe an i.i.d. sample $(K_j^R, K_j, X_j)_{j=1,...,J}$, such that, for each $x \in \{1, \ldots, \overline{X}\}$, the distribution of (K^R, K, ρ) conditional on $\{X = x\}$ satisfies the same conditions as $P^{(K^R, K, \rho)}$ in Assumption 3.1. Thus, we impose the core binomial restriction of Equation (3.1) type by type, and authorize $P^{(K,\rho)|X=x}$ to vary in x; the idea being precisely to take into account the different patterns of polarization according to types. The analysis is identical to the unconditional one but done separately for each type, that is, each subset of options. For any $x \in \{1, \ldots, \overline{X}\}$, we define

$$\pi_x := 1 - \frac{\mathbb{E}[K\rho(1-\rho) \mid X=x]}{2\mathbb{E}[K \mid X=x] q_x(1-q_x)}, \qquad (3.11)$$

where $q_x := \mathbb{E}[K\rho | X = x]/\mathbb{E}[K | X = x]$ is assumed to be in (0, 1). It is the same definition as our index π in Equation (3.2) restricted to options of type X = x. Identification proceeds as before. Likewise, estimation is done as that of π (see below) on the subsample of options $\{j \in \{1, \ldots, J\} : X_j = x\}$. The within-indices $\{\pi_x\}_{x=1,\ldots,\overline{X}}$ are often of interest on their own. We can also consider an aggregated conditional measure, with weights equal to the fraction of each option type,

$$\pi_{\text{conditional}(X)} := \sum_{x=1}^{\overline{X}} \mathbb{P}(X=x) \,\pi_x. \tag{3.12}$$

Individuals' characteristics Alternatively, covariates can be defined at the level of the individuals making choices or, more broadly, at the level of the choices themselves.¹³ The extension is similar. We consider a categorical variable $Z \in \{1, \ldots, \overline{Z}\}, \overline{Z} \geq 2$, that indicates the type of each choice. It is specified at choice-level, although it can inherit from the features of individuals making choices. In our application, it represents the type of speakers, defined as interactions of personal covariates, such as gender, constituency, and chamber (House of Representatives versus Senate).

Again, we consider polarization within each type. For any type $z \in \{1, \ldots, \overline{Z}\}$ and word (option) j, $K_{j,z}^R$ (respectively $K_{j,z}^D$) denotes the number of occurrences (choices) pronounced by Republicans (resp. Democrats) of type z. $K_{j,z} := K_{j,z}^R + K_{j,z}^D$ is the number of occurrences pronounced by speakers of type z. We thus have $K_j^R = \sum_{z=1}^{\overline{Z}} K_{j,z}^R$, and the same equality holds for the variables K^D and K. We define $\rho_{j,z}$ as the probability, conditional on word j being pronounced by a speaker of type z, that the speaker in question is Republican. We assume our DGP assumption holds within each type.

Assumption 3.2 (DGP – Individuals' covariates). We have a sample made of J distinct options and, for each type $z \in \{1, \ldots, \overline{Z}\}$, we observe $(K_{j,z}^R, K_{j,z})_{j=1,\ldots,J}$ i.i.d. random variables such that $(K_{j,z}^R, K_{j,z}, \rho_{j,z})$ has the same distribution as (K_z^R, K_z, ρ_z) , which satisfies: $\mathbb{E}[K_z] > \mathbb{E}[K_z \rho_z] > 0$, and

$$K_z^R | K_z, \rho_z \sim \text{Binomial}(K_z, \rho_z).$$
 (3.13)

¹³To illustrate the distinction, imagine wondering whether the level of partisanship is higher in the morning than in the afternoon, before a lunchtime break where congresspeople might have informal discussions to clarify their positions and, perhaps, reach compromises. The dummy that indicates whether an occurrence (a choice) is pronounced before or after lunch is defined at the level of choices, regardless of speaker identities. The formalization encompasses the two cases, speakers' covariates or characteristics defined at the level of occurrences alike.

We conduct the analysis type by type defining within-type indices by

$$\pi_z := 1 - \frac{\mathbb{E}[K_z \rho_z (1 - \rho_z)]}{2\mathbb{E}[K_z] q_z (1 - q_z)} , \qquad (3.14)$$

where $q_z := \mathbb{E}[K_z \rho_z]/\mathbb{E}[K_z]$ is assumed to be in (0, 1), for each $z \in \{1, \ldots, \overline{Z}\}$. Using the variables K_z^R, K_z^D, K_z , instead of K^R, K^D, K , in Theorem 3.2 and Section 3.3.2 gives corresponding identification results for π_z . A weighted mean of the $\{\pi_z\}_{z=1,\ldots,\overline{Z}}$ provides an aggregated conditional index of polarization.

The definition of Equation (3.14) is the same as the unconditional index's restricted to occurrences (choices) of type Z = z, hence a noteworthy difference with the previous options' covariates setting. With options' covariates, we perform the analysis for each subset of options as if we had separate samples. On the contrary, with individuals' characteristics, we keep the same choice set but successively study occurrences pronounced by the distinct types. As a consequence, conditioning on individual types makes the small-unit issue more acute: the number of times a word is pronounced within a given type, say, by women senators from the West Coast, is smaller than for the full sample. A high \overline{Z} relative to the number of observed choices entails low values for the variables $\{K_z\}_{z=1,...,\overline{Z}}$, possibly with a large fraction of single occurrence. The identification interval for π_z is then likely to be broad.

Estimation and inference for conditional indices Be it with individual or option characteristics, the estimation and inference for the within-indices, π_x or π_z , proceed exactly as for the unconditional index π ; the only modification concerns the sample used. For covariates defined at the level of option, we work with the subsample of phrases (options) $\{j \in \{1, \ldots, J\} : X_j = x\}$. For individual features, we restrict to the occurrences (choices) of type Z = z: $(K_{j,z}^R, K_{j,z})_{j=1,\ldots,J}$. Then, we estimate the aggregated conditional indices, or their identified sets in general, by plug-in. That is, for any type x of option, the quantity $\mathbb{P}(X = x)$ is estimated by the empirical counterpart $J^{-1} \sum_{j=1}^{J} \mathbb{1}\{X_j = x\}$ and, combined with identified sets or estimates for $\{\pi_x\}_{x=1,\ldots,\overline{X}}$, the injection in Equation (3.12) gives an estimate of the identified set or a point-estimate of $\pi_{\text{conditional}(X)}$. The case of covariates defined at the level of occurrences or choices is identical. For the weights, the estimator $\sum_{j=1}^{J} K_{j,z} / \sum_{j=1}^{J} K_j$ estimates the probability that a choice is of type Z = z, for each $z \in \{1, \ldots, \overline{Z}\}$. As regards inference, the aggregated indices are convex combinations of the within-indices whose weights are population moments. As for the unconditional index, we can then follow the previous construction to obtain confidence intervals with asymptotic guarantees.

3.4.2 Test of the binomial assumption

In addition to i.i.d. sampling, our main assumption is the binomial distribution for K^R conditional on K and ρ . That assumption yields a mapping between the distribution of K^R , which is identified, and some features of the distribution of ρ , which is unobserved. As for now, consider K is fixed and known or, equivalently, the reasoning is conditional on K. For any $k \in \mathbb{N}^*$, the binomial assumption gives the identification of the first k moments of the distribution $\mathbb{P}^{\rho|K=k}$. Besides, a vector of moments needs to satisfy some restrictions. For instance, the non-negativity of variance implies that the second moment cannot be lower than the square of the first moment. Formally, it has to belong to the moment space.

The intuition of the test, developed for a fixed K in D'Haultfœuille and Rathelot (2017), is the following. The binomial assumption allows to identify the first moments of $P^{\rho|K=k}$ and, on the other hand, any vector of moments has to belong to the moment space. Therefore, the data can contradict the binomial assumption when the estimated vector of moments is too far, taking into account sampling uncertainty, from the moment space.

3.5 Application to speech polarization in the U.S. Congress (1873-2016)

We apply our method to revisit Gentzkow et al. (2019) on speech partial sanship between Republican and Democrats congresspeople from 1873 to 2016. Following GST, we consider bigrams as the unit of analysis.¹⁴ We show how the identification set and the extrapolated point-estimate of the unconditional polarization index vary over time and how they differ from the naive estimate. Then, we discuss how the choice of the dictionary affects the results and conclude with recommendations on how to choose the dictionary in practical cases. The third subsection presents two tests assessing the reliability of our approach.

3.5.1 The evolution of speech polarization over time

Figure 3.1 presents the evolution of speech polarization over time. In this part of the analysis, we use our preferred choice for the dictionary, which we discuss in the next subsection. The blue curve represents the naive estimate, which ignores the small-sample issue. The area in red delimited by the plain red lines is the confidence interval at 95%, while the red dots (very close to the red lines) are the estimated upper and lower identification bounds. In this application with large J (above one million for our preferred dictionary in each session), the statistical uncertainty is small, and the uncertainty regarding the target parameter π mainly comes from partial identification. Indeed, as explained above, the identification set is not reduced to a point because of the existence of words that are said only once. Our results point to the fact that these words make the identification issue by making some assumptions on the contribution of one-occurrence words to polarization, is displayed in orange.

We draw several conclusions from the results. From a methodological perspective, we note that the naive index is far above the extrapolated index and is even above the upper identification bound. This demonstrates that it is impossible to draw conclusions about the magnitude of polarization from the naive approach, which over-estimates polarization to a large extent. The naive estimate is also wrong in evolution: strongly decreasing until 1975 and increasing afterward while the extrapolated estimate displays a U-shape with the highest levels of polarization in recent years. The explanation is that the bias of the naive estimate has decreased over time because the volume of speech has increased, mitigating small-unit issues (see Figure 3.12 in Appendix 3.C.2).

From a political science perspective, our conclusions are similar to Gentzkow et al. (2019) in the sense that, like them, we document an increase in speech polarization since

 $^{^{14}}$ Appendix 3.C.1 explains the operations that transform the raw transcripts of speeches into the counts of bigrams used in our analysis and shows some examples that compare the raw speech and the final counts of valid bigrams.

the sixties. The level observed since the 2010s is the highest ever reached over the entire period (1873-2016). However, our estimate differs from there in a few aspects. First, contrary to their preferred penalized estimate (Panel B of their Figure 2), we document a fair amount of speech polarization in the very beginning of the 20th century and, to a lower extent, in the late twenties. While current levels are higher, the level reached in 1908 is not very far off. The overall pattern since 1900 is thus rather a U-shape than a continuous increase.

Figure 3.1: Evolution of polarization over time: identified set, extrapolated point-estimate and naive estimate of the polarization index π for our preferred specification regarding dictionary choices.



Note: Naive estimate (blue line), identified set (red area delimited by red dotted lines), 95% confidence interval for the polarization index (plain red lines), extrapolated estimate (orange plain line) and its associated 95% confidence interval (orange dashed lines) (remark: this confidence interval assumes the model $f(\cdot, \theta_0)$ underlying the extrapolation is correct; in other words, it only accounts for statistical uncertainty, but not model uncertainty). Each point corresponds to a Congressional session. The parameters used for the extrapolation are $\overline{k} = 8$ (max number of occurrences) and r = 3 (polynomial degree). We do not include covariates in this analysis.

Sample: dictionary of bigrams after spelling corrections and exclusion of invalid words but without exclusions based on frequency.

3.5.2 Dictionaries, processing operations, and small-unit bias

Before any analysis, it is necessary to choose the dictionary of bigrams that are kept in the data. In most papers, this data-processing step is not salient. We show here how results may vary as a function of the choices made.

The main objectives of this data-processing step are (i) to keep only meaningful bigrams, (ii) to ease computational burden. This step usually consists of several stages:

- removing words that belong to a given neutral vocabulary (e.g., procedural phrases, conventional stop-words, Congress people's names);
- keeping bigrams that are pronounced sufficiently frequently by a sufficient number of different speakers and in a sufficient number of distinct sessions.

Hereafter, we use the list of words provided by GST to remove the neutral vocabulary and focus our analysis on the second stage, which aims at getting rid of rare bigrams. As our objective is to assess polarization, rare bigrams should be removed when they correspond to typos or mistakes, but not necessarily in other cases. Indeed, they might precisely reflect marginal bigrams used by few speakers but revealing differences between party languages.

In Figure 3.2, we show the share of bigrams that are only pronounced once in each session according to the data-processing choices. The top curve shows that share before any processing, above 80% in the first sessions and around 65% since the fifties. Spelling correction helps, but the share of bigrams pronounced only once remains around 60%. Finally, when we follow the same processing as in GST, the share of bigrams pronounced only once is much lower, between 10% and 35% depending on the session. The main reason is that their data processing includes explicit restrictions on bigram frequencies: a bigram is kept in the dictionary if it is pronounced a minimum number of times, by a minimum number of speakers, in a minimum number of sessions. While these restrictions make sense, they are ad hoc by nature, and practitioners hardly control whether the bigrams that are removed are meaningful or not. Interestingly, all these curves exhibit a decreasing pattern until the nineties and are either constant or increasing at the end of the period. Note that because the number of bigrams per session increases over time, the share of one-occurrence bigrams displays a stronger decrease over time with GST processing.





Note: Share of bigrams that only appear once in a session, for each session, and for three different choices of the dictionary. The top green curve corresponds to the raw dictionary (after removing neutral words, but before any spelling correction or frequency restriction). The blue curve corresponds to the same one after spelling correction, before any restriction based on bigram frequency. The bottom red one corresponds to the dictionary obtained after the data processing described in GST.

Figure 3.3 shows the identified interval and the extrapolated estimate for each of the three sets of data-processing choices. When we do not restrict bigrams based on their frequency, the number of once-pronounced bigrams is large, resulting in broad identified intervals. The blue set is slightly narrower than the green one due to the exclusion of typos

and spelling mistakes but remains large. The red set, which corresponds to the restrictions made in GST, is much narrower. This is logical since Theorem 3.2 shows the length of the identification interval is proportional to the fraction of one-occurrence bigrams. Before 1945, the red identification set overlaps largely with the other two. However, after 1960, it is mostly distinct. Looking at the extrapolated estimate, we see that the red curve is overall smoother than the green and the blue ones. The changes that happen at the beginning of the 20th century are less abrupt in the red curve, and the increase at the end of the period of analysis looks more sudden.

The lower the share of once-pronounced bigrams, the tighter the identification set. At first sight, this is a good thing as a tighter set means more precise conclusions. However, data-processing choices may not be neutral to the quantitative or even the qualitative conclusions drawn from the data. ad hoc restrictions might take the estimate out of the identification set obtained without restrictions and conceal part of the phenomenon under study. Because one needs to make choices to get informative results, we argue that the assumptions leading to the extrapolated estimation of the index might be more acceptable and easier to interpret than ad hoc dictionary restrictions based on bigram frequency. Furthermore, these assumptions can be tested to some extent (see next Subsection 3.5.3).

In conclusion, we would recommend first eliminating typos and mistakes with spelling corrections and using a dictionary to eliminate procedural or neutral words that make sense in the case under study. Second, instead of ad hoc data-processing restrictions based on bigram frequency, we would recommend using the extrapolated estimator of the index.

Figure 3.3: Identified set, extrapolated point-estimate and related confidence intervals for the polarization index π under different choices of the dictionary.



Note: Identified intervals (shaded areas), extrapolated point-estimates (plain lines) and their 95% confidence intervals (dashed lines) for different choices of the dictionary: the raw dictionary without spelling correction or frequency restriction (green); the same raw dictionary with spelling correction (blue); the dictionary obtained with frequency restriction following GST data-processing (red). The parameters used for the extrapolation are $\bar{k} = 8$ (max number of occurrences) and r = 3 (polynomial degree). We do not include covariates in this analysis.

3.5.3 Testing our approach

Our approach relies on several assumptions. Two of them are testable. First, we test the binomial assumption of the model. Second, we perform an over-identification test for the extrapolated index.

Figure 3.4 shows the p-value corresponding to the test of the binomial assumption for each session, taking a maximum number of occurrences for a bigram equal to 40 (99% of the bigrams are pronounced at most 40 times in a given session). Overall, we cannot reject the binomial distribution assumption: only session 81 (years 1949-1950) has a p-value lower than 10% (equal to 1%).



Figure 3.4: p-values for the test of the binomial assumption.

Note: For each session, we report the p-value of the test of the null hypothesis that, conditional on K and ρ , K^R follows a Binomial distribution. The p-values are obtained by bootstrap (200 repetitions). The analysis is limited to bigrams $\{j \in \mathcal{J} : K_j \leq 40\}$, that is, bigrams pronounced at most 40 times by session. Such bigrams represent 99% of the set of bigrams. In terms of occurrences, they account for around 70 to 80% of all occurrences pronounced on a given session. This restriction enables to reduce the computational cost of the test. We acknowledge that the results for sessions 43, 111, 112, 113, and 114 are (temporary) missing.

Figure 3.5 shows the instances where the over-identification test for the extrapolated index is rejected at 1%. More precisely, the null hypothesis is that $k \mapsto m(k)$ is a polynomial of order 1, 2, or 3 for k between 2 and 8. The analysis reveals that linear extrapolations are rejected very often. Quadratic extrapolations are also rejected often, especially after the 1960s. Polynomial of order 3, however, are only rejected four times over the whole period with our preferred specification regarding dictionary choices. This result is even more reassuring for the extrapolation approach that the large number J of bigrams makes it relatively easier to reject even small deviations from the null hypothesis. It explains our choice for the extrapolated estimator in Figures 3.3 and 3.1: cubic extrapolation of m(1)from $(m(2), \ldots, m(8))$.

Besides, the analysis reveals that the hypothesis of quadratic or cubic polynomials is more frequently rejected when using the dictionary choice of GST with frequency restrictions compared to our favored data processing. On the other hand, linear extrapolation (r = 1) is instead less rejected for GST dictionary. We have no precise idea to discuss whether those results are mere coincidence or whether frequency restrictions might impact in some way the regularity of the function $k \mapsto m(k)$.



Figure 3.5: Rejections of polynomial assumptions for the extrapolation at 1%.

Note: For each session, polynomial order $r \in \{1, 2, 3\}$, and dictionary choice (ours or GST), the presence of a point means that the over-identification test from CMD of the null hypothesis that $\{2, \ldots, \overline{k} = 8\} \ni k \mapsto m(k)$ is a polynomial of order r is rejected at 1%. "Own bigrams" (orange) corresponds to our preferred dictionary, namely with spelling correction and without restrictions based on bigram frequency; "GST bigrams" (green) corresponds to the dictionary used by GST with frequency restrictions.

3.6 Conclusion

This article proposes a statistical framework and a related methodology to (partially) identify and estimate a polarization index, especially in settings of large choice sets when we observe few choices per option. In such settings, the segregation literature has taught us that naive natural approaches suffer from a small-unit (small-option with our terminology) bias resulting in over-estimation of the level of polarization and impossibility of reliable comparisons over time or across contexts.

Compared to another method proposed in Gentzkow et al. (2019), the pros of our approach consists in, first, formal identification and inference results under an asymptotic in the number J of options tending to infinity, hence consistent with the setting of large choice sets; second, simple and computationally light estimators and confidence intervals. On the other hand, the index of interest is generally partially identified due to a positive fraction of options chosen only once. Depending on the context, that fraction may be substantial. For text analysis, our application to U.S. Congress transcripts evidences that this fraction can lead to uninformative identification interval absent ad hoc selections of the dictionary based on frequency restrictions. In that sense, our partial identification result warns against the possibility that data-processing choices are not neutral to the quantitative, if not qualitative, conclusions on speech polarization. To address this issue, we propose an extrapolated point-estimator.

Another interesting feature of our methodology is that it is, to some extent, testable. We can test the binomial assumption involved in our assumed data-generating process and the model used to construct our extrapolated estimator. Overall, we do not reject those two assumptions in our application to speech polarization in the U.S. Congress.

Nonetheless, our method also has limitations compared to GST methodology and to study text data. The rest of this conclusion focuses on some of them and contemplates possible remedies for future work. Conditional analysis and composition-invariance The first drawback of our approach is its difficulty to perform conditional analysis with speaker-level (or choice-level in the abstract terminology) covariates. As explained in Section 3.4.1, it compounds the small-unit issue. The theory has no real problem with that. However, in practice, even with a few covariates (e.g., gender, chamber, and geographical regions for our application), the number \overline{Z} of possible combinations (or types) is too large for giving informative identified intervals.

A candidate solution would be to use the extrapolated point-estimator. Yet, it appears problematic due to (the lack of) an axiomatic property of the index π we consider. As detailed in Section 3.B.1, π is connected with the Coworker index of segregation. This index is known to fail the composition-invariance property (Frankel and Volij, 2011). Consequently, all else equal, the farther the share of Republican speech (the share of choices made by the minority or reference group in general contexts) from one half, the higher the index. There have been debates about the desirability of composition invariance for segregation indices. When concerned with conditional aggregated indices such as $\pi_{\text{conditional}(X)}$ in Equation (3.12), it seems to us that composition invariance is absolutely necessary. Indeed, in conditional analyses, we are interested in covariates that are good predictors of group membership by construction, whether Republican or Democrat in the application. Therefore, within each different type $z \in \{1, \ldots, \overline{Z}\}$, the proportions of minority and majority groups tend to depart from 1/2 - 1/2. If the index is not composition invariant, the within-type indices π_z will thus mechanically be larger, resulting artificially in a larger aggregated conditional index.

The generalization of the index to allow for unequal group proportions, from π'' to π , considerably weakens the impact of group proportions and brings π closer to composition invariance (see the differences between Figure 3.6 and 3.7 in Appendix 3.B.1)¹⁵ That being said, it remains that π is globally composition variant, that is, sensitive to unequal group proportions. This cast doubts on the reliability of conditional analyses. This is why we restrict to unconditional analyses in Section 3.5.

A possibility to address the problem would be to define another polarization index that is composition invariant, either from classical segregation indices or a new one, possibly more specific to the case of speech polarization.

Specific challenges of text data A second concern relates to the specificity of text data compared to other types of data that might raise concerns about the plausibility of the *i.i.d.* modeling and the definition of the index π as a functional of a joint distribution of (K, ρ) .

A possibility to highlight the specificity of text data is to use array notation. For a corpus of text with a vocabulary or dictionary of size J (be it bigrams, N-grams, embedded words, narratives, etc.), we observe for each $j \in \{1, \ldots, J\}$ the number $K_{j,J}$ of occurrences of the bigram j. Imagine that we observe a larger corpus of text (say, for instance, instead of doing the analysis by Congress session, we perform it by presidential mandates, pooling sessions together). Given Herdan's law (Remark 3.1), both the size J of the dictionary and the total number of occurrences $n = \sum_{j=1}^{J} K_{j,J}$ increase. But it is also sensible that, for each j, the distribution of $K_{j,J}$ also changes by taking larger values since we observe more text. In other words, in the set-up of text data, the number of occurrences K by bigrams (the same would hold for the number of occurrences pronounced by a given party,

¹⁵In addition, points (ii) and (iii) of Theorem 3.1 state that π is *locally* composition invariant in the two extreme cases of null polarization ($\rho \sim \text{Dirac}$) and full polarization ($\rho \sim \text{Bernoulli}$).

 K^R and K^D) are probably better represented by an array

$$(K_{j,J})_{j=1,\dots,J} \stackrel{\text{i.i.d.}}{\sim} K_J$$
 (expressing the dependence in J) or
 $(K_{j,n})_{j=1,\dots,J_n} \stackrel{\text{i.i.d.}}{\sim} K_n$ (expressing the dependence in n , with $J = J(n) = J_n$)

than by a sequence

$$(K_j)_{j=1,\dots,J} \stackrel{\text{i.i.d.}}{\sim} K$$
 (with a distribution of K independent of J or n).

Like GST, this article focuses on a methodological perspective, with the case of speech polarization seen as a particular instance of the more general problem of quantifying polarization in large choice sets. Yet, that distinction between sequence-type modeling and array-type might question the credibility of a unified framework. For applications related to segregation (residential, school, occupational), the sequence modeling seems sensible: if, for instance, the unit/option is a firm or workplace (K is thus the number of individuals who work there and K^R the number of those individuals who are member of a given minority or reference group), more data means observing more workplaces but assuming a common distribution for K and K^R independently of the number J of firms sampled makes sense. It is less arguable for text data.

Recall that the polarization index writes

$$\pi = g(\mathbf{P}^{(K,\rho)}) = 1 - \frac{1}{2q(1-q)} \mathbb{E}\left[\frac{K}{\mathbb{E}[K]}\rho(1-\rho)\right] = 1 - \frac{\mathbb{E}[K\rho(1-\rho)]}{2\mathbb{E}[K\rho]\left(1 - \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}\right)}$$

Consequently, except particular cases of independence between K and ρ (see Section 3.B.7), a distribution of K that depends on J or n entails that the parameter π also depends on J or n. In that sense, it somewhat gets us back to the initial small-unit bias.

Yet, the picture is not so bleak for our methodology and the results presented in the paper. Indeed, as explained in detail in Section 3.D, it is possible to connect the modeling where an observation corresponds to one bigram/option (Assumption 3.1) with the modeling where an observation corresponds to one occurrence/choice (Assumption 3.3). In the latter framework, each new occurrence is pronounced by a Republican with probability q, in this case, the bigram j is pronounced with probability $p_j^{\rm R}$ (respectively with probability $p_j^{\rm D}$ if the occurrence is pronounced by a Democrat). Let define $\rho_j = qp_j^{\rm R}/(qp_j^{\rm R} + (1-q)p_j^{\rm D})$. We consider an infinite dictionary/set of options \mathcal{J} . We can define a polarization index relying solely on q and the primitive probabilities $(p_i^{\rm R}, p_j^{\rm D})_{j \in \mathcal{J}}$ by¹⁶

$$\pi_{J_{\infty}} := 1 + \frac{1}{4q(1-q)} \left[\sum_{j=1}^{\infty} q p_j^{\mathrm{R}} \rho_j + (1-q) p_j^{\mathrm{D}} (1-\rho_j) - 1 \right], \qquad (3.15)$$

and thus independent of the length n of the text corpus (number of observed choices). The interpretation of $\pi_{J_{\infty}}$ as a polarization index is similar to the index π , as suggested by this alternative expression

$$\pi_{J_{\infty}} = 1 - \frac{1}{2q(1-q)} \left[\sum_{j=1}^{\infty} \bar{p}_j \rho_j (1-\rho_j) \right]$$

¹⁶Equation (3.15) is the generalization of index π_J defined in Equation (3.52) in Section 3.D.

where $\bar{p}_j = qp_j^{\rm R} + (1-q)p_j^{\rm D}$ is the weighted average probability that word j occurs. It happens that, in this alternative framework, $\hat{\pi}$ and $\hat{\pi}$ (with the very same expressions as in Equation (3.5) are *also* the estimators of bounds on $\pi_{J_{\infty}}$. This gives a second justification of our bounds in a DGP written at the occurrence level.

Finally, for text data, the last interrogation relates to the independence part of the *i.i.d.* modeling. Remark that the problem concerns both kinds of DGP, either for independence stated at the level of bigrams/items of a dictionary (as in the main body of this article) or at the level of occurrences (as in GST modeling and Section 3.D). In fact, neither bigram counts nor separate occurrences are the native format of the data: texts come in sentences and speeches (Figures 3.9, 3.10 and 3.11 show some examples of speeches from the U.S. Congress debates transcripts).

In this article, following a large part of the literature, we adopt a "bag-of-words" approach that abstracts from the syntactic and grammatical structure and considers words as independents of each other.¹⁷ Yet, that approach might be too coarse, and possible dependence between occurrences within speeches may matter.

We perform analyses on subsamples of the data to assess the importance of such dependence. More precisely, we compare two subsampling schemes: (i) of speeches and (ii) of occurrences stratified by speeches. In the first one, a subsample of the data is obtained by drawing a fraction of the speeches. Therefore, this scheme preserves the speech structure and possible dependencies across occurrences within a speech. In the second one, a subsample of the data is obtained by drawing a fraction of occurrences from each speech (each speech being reduced to a set of occurrences), thus breaking speech and sentence structure.¹⁸ If there is correlation across occurrences within a speech (for instance, a positive correlation if speakers tend to repeat the words they use), that dependence is kept in scheme (i) but reduced in scheme (ii). We do find a difference between the two schemes regarding the proportion of bigrams pronounced once (see Figure 3.16 in the supplements to the application), which suggests the existence of a positive correlation across occurrences within a speech. We leave for future work the study of the importance and impact of such correlation and whether it implies leaving the *i.i.d.* framework or can be accounted for with some speaker- or speech-level unobserved heterogeneity. For the application to U.S. Congress debates, it is worth noting that the number of speeches and their average length vary over time (see Figure 3.15), hence inducing changes in the strength of the possible correlation.

 $^{^{17}}N$ -gram models, with N>1, are a way to nonetheless conserve some of that structure by considering N adjacent words.

 $^{^{18}}$ Section 3.C.2 presents a toy example to illustrate the difference between the two subsampling schemes.

References

- Allen, R., Burgess, S., Davidson, R. and Windmeijer, F. (2015), 'More reliable inference for the dissimilarity index of segregation', *The Econometrics Journal* 18(1), 40–66.
- Berry, S., Levinsohn, J. and Pakes, A. (1995), 'Automobile prices in market equilibrium', Econometrica: Journal of the Econometric Society pp. 841–890.
- Carrington, W. J. and Troske, K. R. (1997), 'On measuring segregation in samples with small units', *Journal of Business & Economic Statistics* 15(4), 402–409.
- Cortese, C. F., Falk, R. F. and Cohen, J. K. (1976), 'Further considerations on the methodological analysis of segregation indices', *American Sociological Review* pp. 630– 637.
- D'Haultfœuille, X. and Rathelot, R. (2017), 'Measuring segregation on small units: A partial identification analysis', *Quantitative Economics* 8(1), 39–73.
- Duncan, O. D. and Duncan, B. (1955), 'A methodological analysis of segregation indexes', American Sociological Review 20(2), 210–217.
- Frankel, D. M. and Volij, O. (2011), 'Measuring school segregation', Journal of Economic Theory 146(1), 1–38.
- Gentzkow, M., Shapiro, J. M. and Taddy, M. (2019), 'Measuring group differences in highdimensional choices: Method and application to congressional speech', *Econometrica* 87(4), 1307–1340.
- Hausman, J. A. (1978), 'Specification tests in econometrics', *Econometrica: Journal of the* econometric society pp. 1251–1271.
- Heaps, H. S. (1978), Information retrieval, computational and theoretical aspects, Academic Press.
- Herdan, G. (1960), Type-token mathematics: A textbook of mathematical linguistics, Vol. 4, Mouton.
- Imbens, G. W. and Manski, C. F. (2004), 'Confidence intervals for partially identified parameters', *Econometrica* 72(6), 1845–1857.
- James, D. R. and Taeuber, K. E. (1985), 'Measures of segregation', Sociological Methodology 15, 1–32.
- Kasy, M. (2019), 'Uniformity and the delta method', *Journal of Econometric Methods* 8(1).
- Kolesár, M. and Rothe, C. (2018), 'Inference in regression discontinuity designs with a discrete running variable', *American Economic Review* **108**(8), 2277–2304.
- Lee, D. S. and Card, D. (2008), 'Regression discontinuity inference with specification error', *Journal of Econometrics* 142(2), 655–674.
- Manski, C. F. (1988), Analog Estimation Methods in Econometrics: Chapman & Hall/CRC Monographs on Statistics & Applied Probability, Chapman and Hall.

- Massey, D. S. and Denton, N. A. (1988), 'The dimensions of residential segregation', Social forces 67(2), 281–315.
- McFadden, D. (1974), 'The measurement of urban travel demand', Journal of public economics **3**(4), 303–328.
- Rathelot, R. (2012), 'Measuring segregation when units are small: A parametric approach', Journal of Business & Economic Statistics **30**(4), 546–553.
- Reardon, S. F. and Firebaugh, G. (2002), 'Measures of multigroup segregation', Sociological Methodology 32(1), 33–67.
- Romano, J. P. (2004), 'On non-parametric testing, the uniform behaviour of the t-test, and related problems', *Scandinavian Journal of Statistics* **31**(4), 567–584.
- Stoye, J. (2009), 'More on confidence intervals for partially identified parameters', *Econometrica* 77(4), 1299–1315.
- Van der Vaart, A. W. (2000), Asymptotic statistics, Vol. 3, Cambridge university press.
- Winship, C. (1977), 'A revaluation of indexes of residential segregation', *Social Forces* **55**(4), 1058–1066.
- Wooldridge, J. M. (2010), *Econometric analysis of cross section and panel data*, MIT press.

Appendix 3.A Proofs of main theorems

3.A.1 Proof of Theorem **3.1** (definition of the index)

Recall that Assumption 3.1 defines a binomial mixture model, namely the distribution of K^R conditional on K and ρ is binomial, whose primitive object is the joint distribution $\mathbb{P}^{(K,\rho)}$. To begin with, remark that $q = \mathbb{E}[K^R]/\mathbb{E}[K]$. Indeed,

$$\mathbb{E}\left[K^{R}\right] \stackrel{\text{LIE}}{=} \mathbb{E}\left[\mathbb{E}\left(K^{R} \mid K, \rho\right)\right] \stackrel{\text{BA}}{=} \mathbb{E}[K\rho], \qquad (3.16)$$

by definition of the expectation of a Binomial distribution, where LIE stands for the Law of Iterated Expectation, and BA for the binomial assumption as expressed in Equation (3.1).

Upper bound of 1 for π The results are almost immediate. Remember the definition

$$\pi := 1 - \frac{\mathbb{E}[K\rho(1-\rho)]}{2\mathbb{E}[K]\,q(1-q)}.$$

Given the nature of the variables K, ρ , and the definition of the quantity q, K, ρ , $1 - \rho$, and q are non-negative. Hence, the numerator and denominator of the fraction are non-negative too, and $\pi \leq 1$, showing the upper bound of (*i*).

In addition, $\pi = 1$ if and only if (iff, henceforth) $\mathbb{E}[K\rho(1-\rho)] = 0$. $K\rho(1-\rho)$ is a non-negative random variable. Therefore, its expectation is null if and only if the variable is almost surely null. Under our DGP, K cannot be constantly equal to 0 since it would contradict the assumption $\mathbb{E}[K] > 0$. Hence, for the expectation to be null, it is necessary and sufficient that, a.s., $\rho(1-\rho)$ is null. The condition characterizes a random variable whose support is restricted to $\{0, 1\}$, namely a Bernoulli variable. Hence, (*iii*) is proved: $\pi = 1$ if and only if \mathbb{P}^{ρ} is a Bernoulli, whatever its parameter.

Lower bound of 1/2 for π Starting from the definition of π , algebraic manipulations immediately give

$$\pi \ge 1/2 \iff \mathbb{E}[K\rho(1-\rho)] \le \mathbb{E}[K]q(1-q).$$

Using the definition of $q := \mathbb{E}[K\rho]/\mathbb{E}[K]$, simplifications and expectation's linearity (EL) enable to write

$$\pi \ge 1/2 \iff \mathbb{E}[K\rho] - \mathbb{E}\left[K\rho^2\right] \le \mathbb{E}[K] \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]} \left(1 - \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}\right)$$
$$\iff \mathbb{E}\left[K\rho^2\right] \ge \frac{\mathbb{E}[K\rho]^2}{\mathbb{E}[K]}.$$

To show the veracity of the latest inequality, we introduce the random variable $a(\rho) := \mathbb{E}[K | \rho] / \mathbb{E}[K]$. By LEI and EL, remark that $\mathbb{E}[a(\rho)] = 1$. Furthermore, $a(\rho) \ge 0$ almost surely. Consequently, we have

$$\mathbb{E}\left(\mathbf{a}(\rho)\left\{\rho - \mathbb{E}[\mathbf{a}(\rho)\rho]\right\}^{2}\right) \ge 0.$$
(3.17)

Expanding the square and using EL, the previous left-hand side equals to

$$\mathbb{E}\left[\mathbf{a}(\rho)\rho^{2}\right] + \mathbb{E}[\mathbf{a}(\rho)] \mathbb{E}[\mathbf{a}(\rho)\rho]^{2} - 2\mathbb{E}[\mathbf{a}(\rho)\rho] \mathbb{E}[\mathbf{a}(\rho)\rho].$$

Using $\mathbb{E}[\mathbf{a}(\rho)] = 1$, Equation (3.17) is equivalent to $\mathbb{E}[\mathbf{a}(\rho)\rho^2] \ge \mathbb{E}[\mathbf{a}(\rho)\rho]^2$. Now, replacing $\mathbf{a}(\rho)$ by its definition, we obtain

$$\mathbb{E}\left[\mathbf{a}(\rho)\rho^{2}\right] \stackrel{\text{def. }\mathbf{a}(\rho)}{=} \mathbb{E}\left(\frac{\mathbb{E}[K \mid \rho]}{\mathbb{E}[K]}\rho^{2}\right) \stackrel{\text{EL}}{=} \frac{1}{\mathbb{E}[K]} \mathbb{E}\left(\mathbb{E}\left[K\rho^{2} \mid \rho\right]\right) \stackrel{\text{LIE}}{=} \frac{\mathbb{E}[K\rho^{2}]}{\mathbb{E}[K]},$$

and likewise, $\mathbb{E}[\mathbf{a}(\rho)\rho] \stackrel{\text{def. }\mathbf{a}(\rho)}{=} \mathbb{E}\left(\frac{\mathbb{E}[K \mid \rho]}{\mathbb{E}[K]}\rho\right) \stackrel{\text{EL}}{=} \frac{1}{\mathbb{E}[K]} \mathbb{E}(\mathbb{E}[K\rho \mid \rho]) \stackrel{\text{LIE}}{=} \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}.$

Thus, Equation (3.17) is equivalent to

$$\frac{\mathbb{E}[K\rho^2]}{\mathbb{E}[K]} \ge \left(\frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}\right)^2 \iff \mathbb{E}[K\rho^2] \ge \frac{\mathbb{E}[K\rho]^2}{\mathbb{E}[K]},$$

which was shown to be a necessary and sufficient condition for $\pi \ge 1/2$. This proves the lower bound in (*i*).

Finally, to demonstrate (*ii*), note that the above reasoning has entirely proceeded by equivalence. As a consequence, π is minimal, equal to 1/2 if and only if Equation (3.17) is an equality. The left-hand side of that equation takes the expectation of a non-negative random variable. Hence, the expectation is null, that is, we have equality, if and only if the variable $a(\rho) \{\rho - \mathbb{E}[a(\rho)\rho]\}^2$ is equal to 0 almost surely.

 $a(\rho) = 0$ entails $\mathbb{E}[K] = 0$. Indeed, by definition of $a(\rho)$, $a(\rho) = 0 \iff \mathbb{E}[K | \rho] = 0$, and, the LEI implies $\mathbb{E}[K] = \mathbb{E}(\mathbb{E}[K | \rho])$. Our DGP assumption excludes that possibility.

Therefore, equality in (3.17) is equivalent to $\rho = \mathbb{E}[a(\rho)\rho]$ almost surely. This means that ρ is a.s. constant and, given a previous computation, equals to $\mathbb{E}[K\rho]/\mathbb{E}[K]$ which is q by definition. This shows (*ii*): $\pi = 1/2$ if and only if \mathbb{P}^{ρ} is a Dirac, and in this case $\rho = q$ almost surely.

3.A.2 Proof of Theorem **3.2** (partial identification)

Given our DGP assumption and the asymptotics in the number of options, the challenge of identification consists in expressing the estimand π as a function of the joint distribution $P^{(K^R, K^D, K)}$. Without further conditions, we are unable to do so. Nonetheless, we obtain an identification interval for π with two sharp bounds.

Identification – isolate $\mathbb{E}[K\rho^2]$ From its definition in Equation 3.2, EL, and $\mathbb{E}[K\rho] = \mathbb{E}[K^R]$, the parameter writes

$$\pi = 1 - \frac{\mathbb{E}\left[K^R\right] - \mathbb{E}\left[K\rho^2\right]}{2\mathbb{E}[K] \frac{\mathbb{E}\left[K^R\right]}{\mathbb{E}[K]} \left(1 - \frac{\mathbb{E}\left[K^R\right]}{\mathbb{E}[K]}\right)}.$$
(3.18)

Therefore, the single term whose identification is to be studied is $\mathbb{E}[K\rho^2]$. To do so, we introduce the function m defined, for any $k \in \mathbb{N}$, by $\mathbf{m}(k) := \mathbb{E}[\rho^2 | K = k]$. It also defines a new random variable $\mathbf{m}(K) := \mathbb{E}[\rho^2 | K]$.

That function will prove useful since we have

$$\mathbb{E}\left[K\rho^{2}\right] \stackrel{\text{LIE}}{=} \mathbb{E}\left[\mathbb{E}\left(K\rho^{2} \mid K\right)\right] \stackrel{\text{EL}}{=} \mathbb{E}\left[K \mathbb{E}\left(\rho^{2} \mid K\right)\right] \stackrel{\text{def. m(·)}}{=} \mathbb{E}[K \mathrm{m}(K)].$$
(3.19)

Hence, the problem is now to identify $\mathbb{E}[K\mathbf{m}(K)]$.

To do so, remark that, as K has support in \mathbb{N} , we can write the equalities

$$Km(K) = Km(K)\mathbb{1}\{K > 1\} + Km(K)\mathbb{1}\{K = 1\}$$

= Km(K)\mathbb{1}\{K > 1\} + m(1)\mathbb{1}\{K = 1\},

since if K is null, so is Km(K). The previous equations state equalities between random variables. A fortiori, their expectations are equal. Hence, taking the expectation and using linearity (remark that m(1) is a number, non-stochastic), and the fact that the expectation of the indicator of any event is the probability of that event, we obtain

$$\mathbb{E}[K\mathbf{m}(K)] = \mathbb{E}[K\mathbf{m}(K)\mathbb{1}\{K > 1\}] + \mathbf{m}(1)\mathbb{P}(K = 1).$$
(3.20)

We study successively the two terms in the right-hand side of that equation.

Identification of the first term in Equation (3.20) Relying on our central binomial assumption, we now show that $\mathbb{E}[Km(K)\mathbb{1}\{K > 1\}]$ is identified. We have

$$\mathbb{E}\left[K^{R}(K^{R}-1) \mid K, \rho\right] \stackrel{\text{EL}}{=} \mathbb{E}\left[(K^{R})^{2} \mid K, \rho\right] - \mathbb{E}\left[K^{R} \mid K, \rho\right]$$
$$\stackrel{\text{BA}}{=} \left[K\rho(1-\rho) + (K\rho)^{2}\right] - (K\rho)$$
$$= -K\rho^{2} + K^{2}\rho^{2} = K(K-1)\rho^{2}, \qquad (3.21)$$

where the second equality uses BA and the expressions of second and first moments from a Binomial (K, ρ) distribution. Again, the previous equality concerns random variables, functions of (K, ρ) . Thus, taking the expectation conditional on K of that equality yields

$$\mathbb{E}\Big(\mathbb{E}\Big[K^R(K^R-1) \mid K, \rho\Big] \mid K\Big) = \mathbb{E}\Big[K(K-1)\rho^2 \mid K\Big]$$
$$\iff \mathbb{E}\Big[K^R(K^R-1) \mid K\Big] = K(K-1)\mathbb{E}\Big[\rho^2 \mid K\Big]$$
$$\iff \mathbb{E}\Big[K^R(K^R-1) \mid K\Big] = K(K-1)\mathbf{m}(K),$$

where the first equivalence use the projection composition property of conditional expectation (the set of functions of K are included in the set of functions of (K, ρ)) and its linearity, while the second only uses the definition of $m(\cdot)$.

The idea next is to divide by K - 1 and take expectation to recover the first term in the right-hand side of Equation (3.20). The indicator and our convention are used to write without the bother of dividing by zero. Remark nonetheless that the issue is deeper than pure writing conventions in so far as the previous equality does not say more than 0 = 0in the event $\{K = 1\}$, and this is precisely why the present argument cannot identify m(1) (see below). Otherwise, multiplying the equality by $\mathbb{1}\{K > 1\}$ and using EL, we have

$$\mathbb{E}\left[K^{R}(K^{R}-1)\mathbb{1}\{K>1\} \mid K\right] = K(K-1)\mathbf{m}(K)\mathbb{1}\{K>1\}.$$
(3.22)

Then using again EL and under the convention that null indicator prevails against dividing by zero (that is, the left-hand side of the next equation is set to 0 if K is not strictly greater than 1), we obtain

$$\mathbb{E}\left[\frac{K^{R}(K^{R}-1)}{K-1}\mathbb{1}\{K>1\} \,|\, K\right] = K\mathbf{m}(K)\mathbb{1}\{K>1\}.$$

Finally, taking the expectation and applying LEI yield

$$\mathbb{E}\left[\frac{K^{R}(K^{R}-1)}{K-1}\mathbb{1}\{K>1\}\right] = \mathbb{E}[K\mathrm{m}(K)\mathbb{1}\{K>1\}].$$
(3.23)

Our DGP identifies the left-hand side of Equation (3.23).

Besides, note that conditioning Equation (3.22) by K = k with $k \ge 2$, we have by EL

$$m(k) = \frac{\mathbb{E}\left[K^{R}(K^{R}-1) \mid K=k\right]}{k(k-1)}.$$
(3.24)

Second term in Equation (3.20) and the source of partial identification With hindsight, Equations (3.18), (3.19), (3.20), and (3.23) show that the only term that remains to be identified in π is m(1). The point is that, absent further assumption, m(1) is not identified, hence causing partial identification for our index. However, we can bound m(1).

In order to do so, note that

$$\rho = \mathbb{E}\Big[K^R \,|\, K = 1, \rho\Big]\,,$$

since, under our binomial assumption, the distribution of K^R conditional on K = 1 and ρ is Binomial with expectation equals to $1 \times \rho = \rho$. Using this equality, we obtain

$$\mathbb{E}[\rho \mid K=1] = \mathbb{E}\left(\mathbb{E}\left[K^R \mid K=1, \rho\right] \mid K=1\right) = \mathbb{E}\left[K^R \mid K=1\right], \quad (3.25)$$

where the first equality uses the previous equation and the second the composition property of conditional expectations.

Upper bound on π **and sharpness** For the upper bound on m(1), observe that, as the support of ρ is [0, 1] (it is a probability), $\rho^2 \leq \rho$ almost surely. Applying expectation conditional on $\{K = 1\}$ preserves that inequality and gives

$$m(1) := \mathbb{E}\left[\rho^2 \,|\, K = 1\right] \le \mathbb{E}[\rho \,|\, K = 1] \stackrel{\text{using (3.25)}}{=} \mathbb{E}\left[K^R \,|\, K = 1\right].$$
(3.26)

In our model, the data identifies $\mathbb{E}[K^R | K = 1]$, thus an upper bound on m(1), and, in the end, an upper bound on π trough Equations (3.18), (3.19), and (3.20).

To show the sharpness of the upper bound, it is necessary and sufficient to exhibit a distribution that (i) is included in our statistical model, (ii) reaches the bound, that is with an equality in Equation (3.26). The latter requirement writes

$$\mathbb{E}\left[\rho^2 \mid K=1\right] = \mathbb{E}[\rho \mid K=1] \iff \mathbb{E}[\rho(1-\rho) \mid K=1] = 0,$$

and, since $\rho(1-\rho)$ is non-negative, it is equivalent to the distribution of ρ conditional on $\{K = 1\}$ being a Bernoulli distribution, possibly degenerate in the sense of a Dirac with mass point in either 0 or 1 (Bernoulli with a parameter equal to 0 or 1). Assumption 3.1 leaves $P^{(K,\rho)}$ unrestricted and the condition $q \in (0,1)$ only rules out that the marginal, unconditional, distribution of ρ is a Dirac in 0 or 1, as it would entail there is in reality one group alone. Hence, our model encompass joint distributions of (K,ρ) such that the conditional distribution of ρ knowing $\{K = 1\}$ is a Bernoulli distribution with any expectation in [0, 1]. Such DGP attain the upper bound on π , proving its sharpness. Besides, remark that a Bernoulli distribution for ρ relates to the highest possible level of polarization, coherent with an upper bound on π . **Lower bound on** π and sharpness The lower bound on m(1) relies on the conditional Jensen inequality. The function $[0,1] \ni z \mapsto z^2$ is convex. Therefore,

$$\mathbf{m}(1) := \mathbb{E}\left[\rho^2 \,|\, K = 1\right] \stackrel{\text{Jensen}}{\geq} \mathbb{E}[\rho \,|\, K = 1]^2 \stackrel{\text{using } (3.25)}{=} \mathbb{E}\left[K^R \,|\, K = 1\right]^2.$$
(3.27)

Again, $\mathbb{E}[K^R | K = 1]$ is identified under our model. A lower bound on m(1), hence on π , is thus identified.

Sharpness of the upper bounds relate to equality in Equation (3.27). Given that the square function $[0,1] \ni z \mapsto z^2$ is strictly convex, Jensen inequality turns out to be an equality if and only if, conditional on $\{K = 1\}$, ρ is almost surely constant. Remember that the unique restriction on $P^{(K,\rho)}$ imposed by our model (except $\mathbb{E}[K] > 0$) is to prevent P^{ρ} being a Dirac in 0 or a Dirac in 1, through the condition $q \in (0,1)$. Thus, joint distributions of (K,ρ) such that the conditional distribution of ρ knowing $\{K = 1\}$ is a Dirac, with any mass point in [0,1] are authorized. Such distribution reach the lower bound on π , which shows that the bound is sharp. Once more, a Dirac distribution for ρ links to the absence of polarization, and logically arises in the lower bound.

Particular cases of point-identification The two previous discussions on sharpness reveal a particular case: if $P^{\rho|K=1}$ is either a Dirac in 0 or a Dirac in 1, the two bounds coincide since $\mathbb{E}[\rho | K = 1] = \mathbb{E}[\rho | K = 1]^2$. In this situation, π is point-identified. However, it would be a rare case in concrete applications.

A situation, likely to be more frequent, in which π is point-identified arises whenever $\mathbb{P}(K = 1) = 0$. Indeed, in this case, the second term in Equation (3.20) is null and the problematic quantity m(1) does not intervene.

Expressions of the bounds as stated in Theorem 3.2 Eventually, the lower and upper bounds on m(1) are

$$\mathbb{E}\left[K^{R} \mid K=1\right]^{2} \le \mathrm{m}(1) \le \mathbb{E}\left[K^{R} \mid K=1\right].$$
(3.28)

To obtain the expression of the bounds given in the Theorem, we begin by simplifying the expression of the index in Equation (3.18). To do so, we use that, by definition of K and EL, $\mathbb{E}[K^D] = \mathbb{E}[K] - \mathbb{E}[K^R]$.

$$\pi \stackrel{\mathrm{Eq.}\,(3.18)}{=} 1 - \frac{\mathbb{E}\left[K^{R}\right] - \mathbb{E}\left[K\rho^{2}\right]}{2\mathbb{E}\left[K\right] \frac{\mathbb{E}\left[K^{R}\right]}{\mathbb{E}\left[K\right]} \left(1 - \frac{\mathbb{E}\left[K^{R}\right]}{\mathbb{E}\left[K\right]}\right)} = 1 - \frac{\mathbb{E}\left[K^{R}\right] - \mathbb{E}\left[K\rho^{2}\right]}{2\mathbb{E}\left[K^{R}\right] \left(\frac{\mathbb{E}\left[K^{D}\right]}{\mathbb{E}\left[K\right]}\right)}$$
$$= 1 - \frac{\mathbb{E}\left[K\right]}{2\mathbb{E}\left[K^{R}\right]\mathbb{E}\left[K^{D}\right]} \left\{\mathbb{E}\left[K^{R}\right] - \mathbb{E}\left[K\rho^{2}\right]\right\}$$
$$= 1 - \frac{\mathbb{E}\left[K\right]}{2\mathbb{E}\left[K^{R}\right]\mathbb{E}\left[K^{D}\right]} \left(\mathbb{E}\left[K^{R}\right] - \mathbb{E}\left[K\rho^{2}\right]\right\}$$
$$= 1 - \frac{\mathbb{E}\left[K\right]}{2\mathbb{E}\left[K^{R}\right]\mathbb{E}\left[K^{D}\right]} \left(\mathbb{E}\left[K^{R}\right] - \mathbb{E}\left[\frac{K^{R}(K^{R}-1)}{K-1}\mathbb{1}\{K>1\}\right] - \mathrm{m}(1)\mathbb{P}(K=1)\right),$$
(3.29)

where the latest equality combines Equations (3.20) and (3.23). It remains to study the ending term, written between parentheses. We insert the bounds from Equation (3.28) and make some simplifications from the definition $K := K^R + K^D$.

Computations for the upper bound By characterization of conditional expectation, we have

$$\mathbb{E}\left[K^{R} \mid K=1\right] = \mathbb{E}\left[K^{R}\mathbb{1}\left\{K=1\right\}\right] / \mathbb{P}(K=1), \qquad (3.30)$$

where, again, simply to be able to write that equality in any case, in the situation where $\mathbb{P}(K=1) = 0$, we make the convention that the previous fraction is equal to 0. By (3.28), the sharp upper bound on $\mathrm{m}(1)\mathbb{P}(K=1)$ is therefore $\mathbb{E}\left[K^{R}\mathbb{1}\left\{K=1\right\}\right]$.

We plug the bound and simplify the resulting expression

$$\left(\cdot \right)_{\text{upper bound}} = \mathbb{E} \left[K^R \right] - \mathbb{E} \left[\frac{K^R (K^R - 1)}{K - 1} \mathbb{1} \{ K > 1 \} \right] - \mathbb{E} \left[K^R \mathbb{1} \{ K = 1 \} \right]$$
$$= \mathbb{E} \left[K^R \mathbb{1} \{ K > 1 \} \right] - \mathbb{E} \left[\frac{K^R (K^R - 1)}{K - 1} \mathbb{1} \{ K > 1 \} \right],$$

where the second equality uses that, since K^R has support in \mathbb{N} , we have the following equality between random variables: $K^R = K^R \mathbb{1}\{K = 1\} + K^R \mathbb{1}\{K > 1\}$ and EL. Now, by EL again and factorization, we have

$$\left(\cdot \right)_{\text{upper bound}} = \mathbb{E} \left[\mathbbm{1} \{K > 1\} \left(K^R - \frac{K^R (K^R - 1)}{K - 1} \right) \right]$$

$$= \mathbb{E} \left[\mathbbm{1} \{K > 1\} \left(\frac{K^R (K - 1) - K^R (K^R - 1)}{K - 1} \right) \right]$$

$$K^R = \underline{K} - K^D \mathbb{E} \left[\mathbbm{1} \{K > 1\} \left(\frac{K^R (K - 1) - K^R (K - K^D - 1)}{K - 1} \right) \right]$$

$$= \mathbb{E} \left[\mathbbm{1} \{K > 1\} \left(\frac{K^R (K - 1) - K^R \left[(K - 1) - K^D \right]}{K - 1} \right) \right]$$

$$= \mathbb{E} \left[\mathbbm{1} \{K > 1\} \frac{K^R K^D}{K - 1} \right].$$

Inserting the result in Equation (3.29) gives the exact upper bound stated in the theorem.

Computations for the lower bound Combining Equations (3.28) and (3.30), the sharp lower bound on $m(1)\mathbb{P}(K=1)$ is $\mathbb{E}[K^R \mathbb{1}\{K=1\}]^2/\mathbb{P}(K=1)$. We plug it in Equation (3.29) to obtain the following lower bound on π :

$$1 - \frac{\mathbb{E}[K]}{2\mathbb{E}[K^R]\mathbb{E}[K^D]} \left(\mathbb{E}\left[K^R\right] - \mathbb{E}\left[\frac{K^R(K^R - 1)}{K - 1}\mathbb{1}\{K > 1\}\right] - \frac{\mathbb{E}\left[K^R\mathbb{1}\{K = 1\}\right]^2}{\mathbb{P}(K = 1)} \right).$$
(3.31)

The term in parentheses remains to be computed. To do so, we can add and subtract $\mathbb{E}[K^R \mathbb{1}\{K=1\}]$. The interest is that we the first two expectations in the parenthesis and the subtracted $\mathbb{E}[K^R \mathbb{1}\{K=1\}]$ precisely form the terms that appear in the previous computation for the upper bound. Therefore, we obtain

$$\left(\cdot\right)_{\text{lower bound}} = \mathbb{E}\left[\mathbbm{1}\{K>1\}\frac{K^{R}K^{D}}{K-1}\right] + \mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right] - \frac{\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{\mathbb{P}(K=1)}.$$

To conclude, we need to compute the difference between the two last terms. We will use the following equality: $\mathbb{1}\{K=1\}(1-K^R)=\mathbb{1}\{K=1\}K^D$. Indeed, either the two

term are null, or K = 1 and then $1 - K^R = K - K^R = K^D$ by definition of K. Using that result and that the expectation of the indicator of any event is the probability of the event, algebraic simplifications give

$$\begin{split} \mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big] &- \frac{\mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big]^2}{\mathbb{P}(K = 1)} = \mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big] \left(1 - \frac{\mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big]}{\mathbb{P}(K = 1)} \right) \\ &= \mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big] \frac{\mathbb{E} \big[\mathbb{1} \{ K = 1 \} \big] - \mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big]}{\mathbb{P}(K = 1)} \\ &= \frac{\mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big] \frac{\mathbb{E} \Big[\mathbb{1} \{ K = 1 \} \Big] \mathbb{E} \Big[K^D \mathbb{1} \{ K = 1 \} \Big]}{\mathbb{P}(K = 1)} \\ &= \frac{\mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big] \mathbb{E} \Big[K^D \mathbb{1} \{ K = 1 \} \Big]}{\mathbb{P}(K = 1)}. \end{split}$$

Therefore, we have

$$\left(\cdot\right)_{\text{lower bound}} = \mathbb{E}\left[\mathbbm{1}\left\{K > 1\right\}\frac{K^{R}K^{D}}{K-1}\right] + \frac{\mathbb{E}\left[K^{R}\mathbbm{1}\left\{K = 1\right\}\right]\mathbb{E}\left[K^{D}\mathbbm{1}\left\{K = 1\right\}\right]}{\mathbb{P}(K=1)},$$

which, plugged in Equation (3.31), yields the lower bound as expressed in the theorem.

Conditional on K > 0 Finally, the last claim follows by remarking that the bounds $\underline{\pi}$ and $\overline{\pi}$ we obtain remain unchanged if we replace all the expectations therein by expectations conditional on K > 0 (Section 3.B.2 in the supplement material to our methodology provides further details).

Appendix 3.B Supplements to the methodology

3.B.1 Connection with the Coworker index of segregation

The literature review in the introduction emphasizes the relation of our method with segregation indices. This appendix shows a formal link between our polarization index π and a common residential segregation index, the Coworker index.

As explained in Section 3.2.2, we follow Gentzkow et al. (2019) as regards the specific form of the index. That specification relates to an index of residential or occupational segregation, whose numerous names include notably the Coworker, the Isolation, or the Normalized Exposure index. Footnote 11, on page 1314, of GST mentions that link, and explains that, in the particular case of "an infinite population with an equal share of Republicans and Democrats", the isolation index is an affine transformation of GST's partisanship index, and they are thus closely related.

An historic definition of the Coworker, in a census framework and using realized proportions, can be found in Equation (3) of James and Taeuber (1985), p6. To address the small-unit issue, we consider a functional of the underlying probabilities ρ instead. In that perspective, Equation (2.1) of D'Haultfœuille and Rathelot (2017), p43, defines the Coworker index as a function of the distribution of ρ .¹⁹ We recall the definition:

$$CW := \frac{\mathbb{E}[\rho^2] - \mathbb{E}[\rho]^2}{\mathbb{E}[\rho] (1 - \mathbb{E}[\rho])} = \frac{\mathbb{V}[\rho]}{\mathbb{E}[\rho] (1 - \mathbb{E}[\rho])}.$$
(3.32)

With a constant and known K, the following computations show that $CW = 2\pi - 1$.

$$\begin{aligned} \pi &:= 1 - \frac{\mathbb{E}[K\rho(1-\rho)]}{2\mathbb{E}[K] \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]} \left(1 - \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}\right)} = 1 - \frac{\mathbb{E}[K\rho(1-\rho)]}{2\mathbb{E}[K\rho] \left(1 - \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}\right)} \\ & \overset{\text{With } K \text{ constant}}{=} 1 - \frac{\mathbb{E}[\rho(1-\rho)]}{2\mathbb{E}[\rho] \left(1 - \mathbb{E}[\rho]\right)}. \end{aligned}$$

Therefore, we have

$$2\pi - 1 \stackrel{\text{EL}}{=} 1 - \frac{\mathbb{E}[\rho] - \mathbb{E}[\rho^2]}{\mathbb{E}[\rho] (1 - \mathbb{E}[\rho])} = \frac{\mathbb{E}[\rho] - \mathbb{E}[\rho]^2 - \mathbb{E}[\rho] + \mathbb{E}[\rho^2]}{\mathbb{E}[\rho] (1 - \mathbb{E}[\rho])}$$
$$= \frac{\mathbb{E}[\rho^2] - \mathbb{E}[\rho]^2}{\mathbb{E}[\rho] (1 - \mathbb{E}[\rho])} =: \text{CW}.$$

Our polarization index is thus connected to a classical segregation index.

In this paper, we take as given the polarization index defined by GST, and adapt it to a statistical framework that yields formal identification and estimation results and is suited to large dimensional choice sets. Consequently, although Theorem 3.1 shows π satisfies the basic requirements to be used as a polarization index, we leave aside the question of the axiomatic properties of π compared to other specifications of polarization indices.

That being said, the link with the Coworker index calls for a short digression about composition invariance. We already alluded to that notion in the Introduction, and this

¹⁹The index π is a function of $P^{(K,\rho)}$. Here, as the objective is only to explicit the link between our polarization index and the Coworker, we restrict to the the case of a constant and known K.

is what motivates the generalization of π'' to unequal group proportions to define π . Intuitively and without formalization, an index is said to satisfy the composition invariance property when it is unaffected by a change in the proportion of the two groups that leaves the differences, the dispersion in the choices made by the two groups identical. To clarify, points (*ii*) and (*iii*) of our Theorem 3.1 shows that our index π satisfies the composition invariance property *locally*. By that affirmation, we mean that in the case of complete polarization, when P^{ρ} is a Bernoulli such that each option is chosen by only one of the two groups, the index π is maximal, equal to 1, whatever the parameter of the Bernoulli distribution, that is whatever the composition of the population. Likewise, in the reverse case of null polarization, $\pi = 1/2$ for any Dirac distribution, whatever its mass point. Nonetheless, to be said invariant to composition, π should satisfy the property *globally*, not only for that two polar cases.





Note: We assume here independence (or sufficiently strong uncorrelation; see Section 3.B.7) between K and ρ . In this case, the index is a function of the first two moments of the distribution of ρ . The vector of those moments needs to satisfy restrictions to belong to the moment space (e.g., non-negative variance), which implies that $\mathbb{E}[\rho]$ cannot be too far from 1/2 for a fixed variance; reversely, for a fixed $\mathbb{E}[\rho]$, $\mathbb{V}[\rho]$ cannot lie anywhere. For a fixed $\mathbb{V}[\rho]$ (expressed as a percentage of the maximum possible variance for $\rho \in [0, 1]$, namely 0.25 attained by a Bernoulli with expectation 0.5), the curve of the corresponding color represents the value of π as a function of $\mathbb{E}[\rho]$ (the share of Republican speech in our application, or the proportion of the minority group in general contexts). In such cases, composition-invariant indices should be sensitive to $\mathbb{V}[\rho]$ only. It is not the case of π in general, although the dependence is mild for $\mathbb{E}[\rho]$ around 1/2 and limited magnitude of polarization (small $\mathbb{V}[\rho]$).

Despite that the formal investigation of such property is beyond the objective of this paper, we mention the point in reaction to the assertion of Gentzkow et al. (2019) that "Ignoring covariates \mathbf{x} , [their] measure satisfies six of these axioms [those of Frankel and Volij (2011)]: Non-triviality, Continuity, Scale Invariance, Symmetry, Composition Invariance, and the School Division Property., GST, p1314. As regards composition invariance, the affine link with the Coworker or Isolation index suggests the contrary given our understanding of the literature. Indeed, although the proofs we are aware of
concern the definition of the coworker index by empirical proportions, instead of underlying probabilities, to quote Frankel and Volij (2011), "It is well known that the [Isolation] index is not Composition Invariant", p12.²⁰ GST's and our index are not formally concerned by those proofs since they are defined directly upon the probabilities. However, absent proof of the contrary, it is plausible that π and the related GST's index are not globally composition invariant.

In fact, in the simplified case where K and ρ are independent (or sufficiently not correlated; see Section 3.B.7), simulation studies show the partial probability index π is not globally composition invariant (Figure 3.6). Nonetheless, we stress that the generalization from π'' to π , allowing for unequal group proportions, weakens the impact of group proportion (compare Figure 3.6 with Figure 3.7).

Figure 3.7: Value of the index π'' (not generalized to unequal group proportions) as a function of $\mathbb{E}[\rho]$ and $\mathbb{V}[\rho]$ when $K \perp \rho$.



Note: See the Note of Figure 3.6 where, instead of π , the graph shows the value of π'' . π'' takes for reference half-half proportions for the two groups and, therefore, is more sensitive to the composition than π .

3.B.2 Restriction to K > 0 without loss of generality

This appendix gives details on the last claim of Theorem 3.2. The short message is the following: the level and measure of polarization, namely the index π and the identified bounds $\underline{\pi}$ and $\overline{\pi}$, are unaffected by irrelevant options, that is with K = 0. Formally, those three quantities, first presented with unconditional expectations, can be rewritten with expectations conditional on $\{K > 0\}$. In other words, they are functions of the conditional joint distribution $P^{(K,\rho)|K>0}$ for the index, respectively $P^{(K^R,K)|K>0}$ for the bounds. As a consequence, an equivalent presentation of our DGP Assumption 3.1 would be to impose K > 0; J being then the number of distinct bigrams pronounced at least one in the data.

²⁰Conceptually, as suggested by its very name, the Coworker or Isolation index seems to depend, by construction, on the overall population of minority. Actually, Massey and Denton (1988) classifies the index as an "exposure" type of segregation indices, see §2 Exposure, p287 of that paper, connecting it to the composition of the population, in contrast with "unevenness" indices.

Expression of the index Our index is defined as a function of the joint distribution of $\mathbb{P}^{(K,\rho)}$ in Equation (3.2). Simply swapping q for its definition $\mathbb{E}[K\rho]/\mathbb{E}[K]$, it writes

$$\begin{aligned} \pi &:= 1 - \frac{\mathbb{E}[K\rho(1-\rho)]}{2\mathbb{E}[K]\frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}\left(1 - \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}\right)} \stackrel{\text{EL}}{=} 1 - \frac{\mathbb{E}[K\rho] - \mathbb{E}\left[K\rho^2\right]}{2\mathbb{E}[K\rho]\frac{\mathbb{E}[K(1-\rho)]}{\mathbb{E}[K]}} \\ &= 1 - \frac{\mathbb{E}[K]}{2\mathbb{E}[K(1-\rho)]} + \frac{\mathbb{E}[K]\mathbb{E}[K\rho^2]}{2\mathbb{E}[K\rho]\mathbb{E}[K(1-\rho)]}. \end{aligned}$$

Now, remark that, as the support of K is N and the support of ρ is [0, 1], we have the following equalities between random variables:

$$\begin{split} K &= K \mathbb{1}\{K > 0\}, \, K\rho = K\rho \mathbb{1}\{K > 0\}, \, K(1-\rho) = K(1-\rho) \\ \mathbb{1}\{K > 0\}, \, K\rho^2 &= K\rho^2 \mathbb{1}\{K > 0\}. \end{split}$$

For each of those variables $K, K\rho, K(1-\rho), K\rho^2$, generically denoted by V, we thus have

$$\mathbb{E}[V] = \mathbb{E}[V\mathbb{1}\{K > 0\}] = \mathbb{E}[V \mid K > 0] \times \mathbb{P}(K > 0), \qquad (3.33)$$

where the first equality uses the previous equation and the second stems from the characterization of conditional expectation.

Since the two fractions involved in the expression of π have the same number of expectations at the numerator and at the denominator, the term $\mathbb{P}(K > 0)$ cancels out when we use the result of Equation (3.33). Therefore, our polarization index is equal to

$$\pi = 1 - \frac{\mathbb{E}[K \mid K > 0]}{2\mathbb{E}[K(1-\rho) \mid K > 0]} + \frac{\mathbb{E}[K \mid K > 0] \mathbb{E}[K\rho^2 \mid K > 0]}{2\mathbb{E}[K\rho \mid K > 0] \mathbb{E}[K(1-\rho) \mid K > 0]}.$$
(3.34)

Thus, it rewrites as a function of the joint distribution of (K, ρ) conditional on $\{K > 0\}$. In that respect, options never chosen (K = 0) are irrelevant.

Expression of the bounds The argumentation is similar for the identification bounds. Under our DGP, namely Assumption 3.1 and $q \in (0, 1)$, Theorem 3.2 defines the bounds as functions of the joint distribution of (K^R, K) :

$$\begin{split} \underline{\pi} &:= 1 - \frac{\mathbb{E}[K] \,\mathbb{E}\left[\mathbbm{1}\{K > 1\} K^R K^D / (K-1)\right]}{2\mathbb{E}[K^R] \mathbb{E}[K^D]} - \frac{\mathbb{E}[K] \,\mathbb{E}\left[K^R \mathbbm{1}\{K = 1\}\right] \mathbb{E}\left[K^D \mathbbm{1}\{K = 1\}\right]}{2\mathbb{E}[K^R] \mathbb{E}[K^D] \mathbb{E}[\mathbbm{1}\{K = 1\}]},\\ \overline{\pi} &= 1 - \frac{\mathbb{E}[K] \,\mathbb{E}\left[\mathbbm{1}\{K > 1\} K^R K^D / (K-1)\right]}{2\mathbb{E}[K^R] \mathbb{E}[K^D]} \end{split}$$

 K, K^R , and K^D have support in \mathbb{N} . Moreover, by definition of $K, K^R > 0 \implies K > 0$, idem for a positive K^D . Therefore, we have these equalities between random variables:

$$K = K \mathbb{1}\{K > 0\}, \ K^R = K^R \mathbb{1}\{K > 0\}, \ K^D = K^D \mathbb{1}\{K > 0\}.$$

Similarly, it is evident that $K > 1 \implies K > 0$, as well as, $K = 1 \implies K > 0$. Hence,

$$\begin{split} &K^R \mathbb{1}\{K=1\} = K^R \mathbb{1}\{K=1\} \mathbb{1}\{K>0\}, \quad K^D \mathbb{1}\{K=1\} = K^D \mathbb{1}\{K=1\} \mathbb{1}\{K>0\}, \\ &\mathbb{1}\{K=1\} = \mathbb{1}\{K=1\} \mathbb{1}\{K>0\}, \quad \frac{K^R K^D \mathbb{1}\{K>1\}}{K-1} = \frac{K^R K^D \mathbb{1}\{K>1\}}{K-1} \mathbb{1}\{K>0\}. \end{split}$$

Therefore, for each random variable V that appears within the unconditional expectations involved in the expressions of $\underline{\pi}$ and $\overline{\pi}$, we have again the equality of Equation (3.33). The fractions that compose the bounds have the same number of expectations at the numerator and at the denominator, and the term $\mathbb{P}(K > 0)$ therefore cancels out when using Equation (3.33) in the definition of the bounds. $\underline{\pi}$ and $\overline{\pi}$ are therefore equal to:

$$\underline{\pi} := 1 - \frac{\mathbb{E}[K \mid K > 0]}{2\mathbb{E}[K^R \mid K > 0] \mathbb{E}[K^D \mid K > 0]} \Big\{ \mathbb{E} \Big[\frac{K^R K^D}{K - 1} \mathbb{1}\{K > 1\} \mid K > 0 \Big] \\ + \frac{\mathbb{E} \Big[K^R \mathbb{1}\{K = 1\} \mid K > 0 \Big] \mathbb{E} \Big[K^D \mathbb{1}\{K = 1\} \mid K > 0 \Big]}{\mathbb{E} \big[\mathbb{1}\{K = 1\} \mid K > 0 \big]} \Big\}, \\ \overline{\pi} := 1 - \frac{\mathbb{E}[K \mid K > 0]}{2\mathbb{E}[K^R \mid K > 0] \mathbb{E}[K^D \mid K > 0]} \mathbb{E} \Big[\frac{K^R K^D}{K - 1} \mathbb{1}\{K > 1\} \mid K > 0 \Big], \quad (3.35)$$

and are thus functions of $P^{(K^R,K)|K>0}$; remember that $K^D = K - K^R$ by definition of K.

In conclusion, options that are not chosen, with K = 0, can be dropped without affecting the index π or the bounds $\underline{\pi}$ and $\overline{\pi}$.

Equivalent DGP assumption with positive K The previous equalities (3.34) and (3.35) allow a final remark about our statistical model. In addition to i.i.d. sampling and the binomial condition, Assumption 3.1 imposes $\mathbb{E}[K] > 0$. Hence, the assumption authorizes K = 0; it only imposes that K is not almost surely null, otherwise there is nothing to think about anyway. Since taking or not taking into account the options with K = 0 does not modify the index nor the bounds, we could instead have chosen the condition $\mathbb{P}(K > 0) = 1$ in Assumption 3.1. That condition is stronger, for it is sufficient to have $\mathbb{E}[K] > 0$, but without loss of generality as regards the definition and partial identification of our index.²¹

Regarding estimation, excluding or not options never chosen in the data, $K_j = 0$, leaves the estimators of the bounds and the extrapolated point-estimator unchanged (see Equations (3.5), (3.9), and (3.10)).

3.B.3 Formal results for estimation and inference

The expression of the identification bounds in Theorem 3.2 suggests simple estimators by the method of moments. Swapping expectations $\mathbb{E}[\cdot]$ for sample means $J^{-1} \sum_{j=1}^{J} \cdot_j$ and simplifying by J, we obtain $\hat{\pi} :=$

$$1 - \frac{\sum_{j=1}^{J} K_{j}}{2\sum_{j=1}^{J} K_{j}^{R} \sum_{j=1}^{J} K_{j}^{D}} \left\{ \sum_{j=1}^{J} \frac{K_{j}^{R} K_{j}^{D}}{K_{j} - 1} \mathbb{1}\{K_{j} > 1\} + \frac{\sum_{j=1}^{J} K_{j}^{R} \mathbb{1}\{K_{j} = 1\} \sum_{j=1}^{J} K_{j}^{D} \mathbb{1}\{K_{j} = 1\}}{\sum_{j=1}^{J} \mathbb{1}\{K_{j} = 1\}} \right\}$$

and $\hat{\pi} := 1 - \frac{\sum_{j=1}^{J} K_{j}}{2\sum_{j=1}^{J} K_{j}^{R} \sum_{j=1}^{J} K_{j}^{D}} \sum_{j=1}^{J} \frac{K_{j}^{R} K_{j}^{D}}{K_{j} - 1} \mathbb{1}\{K_{j} > 1\}.$

For the proper definition of those estimators, we extend the natural convention of Theorem 3.2 to empirical counterparts. First, we set $\mathbb{1}{K_i > 1}/(K_i - 1)$ to 0 if $K_i = 1$.

²¹In the perspective of our asymptotics in J tending to $+\infty$, the condition $\mathbb{E}[K] > 0$ is sufficient to ensure that the sequence of observed $\{K_j\}_{j \in \mathcal{J}}$ cannot be always null, and therefore does give some information on the underlying distribution.

Second, we take $\sum_{j=1}^{J} K_j^R \mathbb{1}\{K_j = 1\} / \sum_{j=1}^{J} \mathbb{1}\{K_j = 1\} = 0$ if $\sum_{j=1}^{J} \mathbb{1}\{K_j = 1\}$ is null. In other words, the last fraction in $\hat{\underline{\pi}}$ vanishes if there is no option with a single occurrence, $K_j = 1$, in the data. In that case, the two estimators coincide and we obtain a point-estimator $\hat{\underline{\pi}} = \hat{\overline{\pi}}$ of our polarization index π .

The computation of $\hat{\underline{\pi}}$ and $\hat{\overline{\pi}}$ is straightforward and numerically light. They display the usual desired properties of method of moments estimators.

Proposition 3.1 (Consistency of $\hat{\pi}$ and $\hat{\pi}$). If Assumption 3.1 holds and $\mathbb{E}[|K|] < +\infty$, then $\hat{\pi}$ and $\hat{\pi}$ are well-defined with probability approaching one as the sample size J grows to infinity, and they consistently estimate the identification bounds: $\hat{\pi} \xrightarrow[J \to +\infty]{a.s.} \pi$, and $\hat{\pi} \xrightarrow[J \to +\infty]{a.s.} \pi$.

The proof consists in applying the Law of Large Numbers and the Continuous Mapping Theorem. Appendix 3.B.4 gives the exact arguments.

Proposition 3.2 (Joint asymptotic normality of $\hat{\pi}$ and $\hat{\pi}$). If Assumption 3.1 holds and $\mathbb{E}[K^2] < +\infty$, then $\hat{\pi}$ and $\hat{\pi}$ are jointly asymptotically normal, namely

$$\sqrt{J}\left[\left(\begin{array}{c} \frac{\widehat{\pi}}{\widehat{\pi}} \end{array}\right) - \left(\begin{array}{c} \frac{\pi}{\pi} \end{array}\right)\right] \xrightarrow[J \to +\infty]{} \mathcal{N}\left(0, \left(\begin{array}{c} \omega & \tau \\ \tau & \overline{\omega} \end{array}\right)\right)$$

with $\underline{\omega}$, τ , and $\overline{\omega}$ numbers that are continuous functions of the expectations of, covariances between, and variances of the variables K, K^R , K^D , $K^R K^D \mathbb{1}\{K = 1\}/(K - 1)$, $K^R \mathbb{1}\{K = 1\}, K^D \mathbb{1}\{K = 1\}, and \mathbb{1}\{K = 1\}.$

Appendix 3.B.5 presents the proof in detail, which relies on applications of the Central Limit Theorem and delta method, and gives the expressions of $\underline{\omega}, \tau$, and $\overline{\omega}$. The expressions are not particularly enlightening and are thus reported there. The important is that, using again the method of moments, we have immediate estimators $\underline{\hat{\omega}}, \hat{\tau}$, and $\overline{\hat{\omega}}$ for the entries of the asymptotic variance-covariance matrix (see Equations (3.41), (3.42), and (3.43) in the appendix). Under the conditions of Proposition 3.2, they are consistent estimators of $\underline{\omega}, \tau$, and $\overline{\omega}$. In practice, this enables to construct a confidence interval (CI) for π with asymptotic guarantees. The CI does concern π , the polarization index of interest, not the identification bounds $\underline{\pi}$ and $\overline{\pi}$ or the identified set $[\underline{\pi}, \overline{\pi}]$. It needs to take into account partial identification of π . To do so, we follow Imbens and Manski (2004) and Stoye (2009).

For $\alpha \in (0, 1)$, we define our CI by

$$\operatorname{CI}_{1-\alpha}^{\pi} := \left[\underline{\widehat{\pi}} - \frac{\operatorname{q}(\alpha)\sqrt{\widehat{\omega}}}{\sqrt{J}}, \overline{\pi} + \frac{\operatorname{q}(\alpha)\sqrt{\widehat{\omega}}}{\sqrt{J}} \right]$$

where $q(\alpha)$ solves

$$\Phi\left(\mathbf{q}(\alpha) + \frac{\sqrt{J}\widehat{\Delta}}{\max\left\{\sqrt{\widehat{\omega}}, \sqrt{\widehat{\omega}}\right\}}\right) - \Phi(-\mathbf{q}(\alpha)) = 1 - \alpha,$$

with $\widehat{\Delta} := \widehat{\pi} - \widehat{\pi}$ and $\Phi(\cdot)$ the cumulative distribution function of $\mathcal{N}(0,1)$ distribution.

Proposition 3.3 (Asymptotic confidence interval for π). If there exist $\varepsilon > 0$, and positive and finite constants B, ω_l , and ω_u such that $\mathbb{E}[|K|^{2+\varepsilon}] < B$, $\omega_l \leq \omega \leq \omega_u$, and $\omega_l \leq \overline{\omega} \leq \omega_u$ for any DGP satisfying Assumption 3.1, then, for any $\alpha \in (0, 1)$,

$$\lim_{J\to+\infty} \mathbb{P}\left(\pi \in CI_{1-\alpha}^{\pi}\right) \ge 1-\alpha.$$

The proof relies on Lemma 3 and Proposition 1 of Stoye (2009). Appendix 3.B.6 provides details and try to give some intuition on the construction of $CI_{1-\alpha}^{\pi}$: for partially identified parameters, the critical value $q(\alpha)$ balances sampling uncertainty with the uncertainty coming from partial identification.

3.B.4 Proof of Proposition 3.1

To make explicit the sample means, the estimators of the bounds rewrite

$$\begin{aligned} \widehat{\pi} &= 1 - \frac{\frac{1}{J} \sum_{j=1}^{J} K_j}{2 \frac{1}{J} \sum_{j=1}^{J} K_j^R \times \frac{1}{J} \sum_{j=1}^{J} K_j^D} \left\{ \frac{1}{J} \sum_{j=1}^{J} \frac{K_j^R K_j^D}{K_j - 1} \mathbb{1}\{K_j > 1\} + \\ & \frac{\frac{1}{J} \sum_{j=1}^{J} K_j^R \mathbb{1}\{K_j = 1\} \times \frac{1}{J} \sum_{j=1}^{J} K_j^D \mathbb{1}\{K_j = 1\}}{\frac{1}{J} \sum_{j=1}^{J} \mathbb{1}\{K_j = 1\}} \right\}, \\ \widehat{\pi} &= 1 - \frac{\frac{1}{J} \sum_{j=1}^{J} K_j}{2 \frac{1}{J} \sum_{j=1}^{J} K_j^R \times \frac{1}{J} \sum_{j=1}^{J} K_j^D} \times \frac{1}{J} \sum_{j=1}^{J} \frac{K_j^R K_j^D}{K_j - 1} \mathbb{1}\{K_j > 1\}. \end{aligned}$$

Definition of the estimators of the bounds Under Assumption 3.1, the restriction $q \in (0, 1)$ implies $\mathbb{E}[K^D] > 0$ and $\mathbb{E}[K^D] > 0$.

Moreover, $K^R \leq K$ and $K^D \leq K$, so that if K has a finite first order moment, $\mathbb{E}[|K|] < +\infty$, so do K^R and K^D . Our DGP assumption also postulates i.i.d. sampling.

Applying the Law of Large Numbers (LLN), the two terms at the denominator in the common fraction of $\hat{\pi}$ and $\hat{\pi}$, $J^{-1} \sum_{j=1}^{J} K_j^R$ and $J^{-1} \sum_{j=1}^{J} K_j^D$, are different from 0 with a probability approaching one as the sample size J grows to infinity. Combined with our convention, a null indicator prevails over a related null denominator, it shows that the two estimators are well-defined with probability approaching one when $J \to +\infty$.

In practice, if it is not the case, namely if $\sum_{j=1}^{J} K_{j}^{R}$ or $\sum_{j=1}^{J} K_{j}^{D}$ is null, the sample at hand contains choices made by only one group, and appears useless to study group differences.

Consistency of the estimators of the bounds To use the Continuous Mapping Theorem (CMT), and later the delta method, we define two real-valued functions for any real numbers $a, b \neq 0, c \neq 0, d, e, f$, and $g \neq 0$,

$$\underline{\mathbf{g}}(a,b,c,d,e,f,g) := 1 - \frac{a}{2bc} \left\{ d + \frac{ef}{g} \right\},$$
(3.36)

$$\overline{\mathbf{g}}(a, b, c, d, e, f, g) := 1 - \frac{a}{2bc} d.$$
 (3.37)

Those functions are continuous. We use the same arguments for both functions although the second one does not make use of the arguments e, f, g because we will consider a multi-valued function $g := (\underline{g}, \overline{g})$ into \mathbb{R}^2 to prove asymptotic normality.

We define the following shortcuts for simplicity

$$a := \mathbb{E}[K], \qquad \qquad \widehat{a} := \frac{1}{J} \sum_{j=1}^{J} K_j,$$
$$b := \mathbb{E}[K^R], \qquad \qquad \widehat{b} := \frac{1}{J} \sum_{j=1}^{J} K_j^R,$$

$$\begin{split} c &:= \mathbb{E} \Big[K^D \Big], & \widehat{c} := \frac{1}{J} \sum_{j=1}^J K_j^D, \\ d &:= \mathbb{E} \Big[\frac{K^R K^D}{K - 1} \mathbb{1} \{ K > 1 \} \Big], & \widehat{d} := \frac{1}{J} \sum_{j=1}^J \frac{K_j^R K_j^D}{K_j - 1} \mathbb{1} \{ K_j > 1 \}, \\ e &:= \mathbb{E} \Big[K^R \mathbb{1} \{ K = 1 \} \Big], & \widehat{e} := \frac{1}{J} \sum_{j=1}^J K_j^R \mathbb{1} \{ K_j = 1 \}, \\ f &:= \mathbb{E} \Big[K^D \mathbb{1} \{ K = 1 \} \Big], & \widehat{f} := \frac{1}{J} \sum_{j=1}^J K_j^D \mathbb{1} \{ K_j = 1 \}, \\ g &:= \mathbb{E} \big[\mathbb{1} \{ K = 1 \} \big], & \widehat{g} := \frac{1}{J} \sum_{j=1}^J \mathbb{1} \{ K_j = 1 \}. \end{split}$$

By construction, we have

$$\underline{\underline{g}}(a, b, c, d, e, f, g) = \underline{\pi}, \qquad \underline{\underline{g}}(\widehat{a}, \widehat{b}, \widehat{c}, d, \widehat{e}, \widehat{f}, \widehat{g}) = \underline{\widehat{\pi}}, \\ \overline{\underline{g}}(a, b, c, d, e, f, g) = \overline{\pi}, \qquad \overline{\underline{g}}(\widehat{a}, \widehat{b}, \widehat{c}, \widehat{d}, \widehat{e}, \widehat{f}, \widehat{g}) = \widehat{\pi}.$$

Before applying the CMT, it remains to verify the conditions of the LLN. Assumption 3.1 posits i.i.d. sampling. Furthermore, we assume $\mathbb{E}[|K|] < +\infty$. That entails a finite first moment for all the relevant variables, namely $K, K^R, K^D, K^R K^D \mathbb{1}\{K > 0\}/(K-1), K^R \mathbb{1}\{K = 1\}, K^D \mathbb{1}\{K = 1\}, \text{ and } \mathbb{1}\{K = 1\}.$

Indeed, K^R and K^D are smaller or equal to K. A fortiori, it is also the case for the variables $K^R \mathbb{1}\{K=1\}$ and $K^D \mathbb{1}\{K=1\}$. The variable $\mathbb{1}\{K=1\}$ is bounded. Finally, we can show that

$$\frac{K^{R}K^{D}}{K-1}\mathbb{1}\{K>1\} \le K.$$
(3.38)

By definition of K, we have $K^R K^D = K^R (K - K^R) \leq K(K - 1)$. To prove the latest inequality, we separate two cases. If K^R is equal to either 0 or K, the left-hand side of the inequality, $K^R (K - K^R)$, is zero and the inequality holds. Otherwise, as the variables in question are integers, we have $1 \leq K^R \leq K - 1$. It gives $K - K^R \leq K - 1 \leq K$. The combination of $K^R \leq K - 1$ and $K - K^R \leq K$ gives $K^R (K - K^R) \leq K(K - 1)$. To obtain the inequality stated in Equation (3.38), note that it holds whenever $K \leq 1$ for the indicator is null, and, for K > 1, dividing the inequality $K^R (K - K^R) \leq K(K - 1)$ by K - 1 > 0 yields the result.

Thus, i.i.d. sampling and the condition $\mathbb{E}[|K|] < +\infty$ allow to apply the strong LLN, which gives, using the previous notations,

$$\forall x \in \{a, b, c, d, e, f, g\}, \ \widehat{x} \xrightarrow[J \to +\infty]{\text{a.s.}} x.$$

The Continuous Mapping Theorem gives the result

$$\underline{\mathbf{g}}(\widehat{a},\widehat{b},\widehat{c},\widehat{d},\widehat{e},\widehat{f},\widehat{g}) = \underline{\widehat{\pi}} \xrightarrow[J \to +\infty]{a.s.} \underline{\mathbf{g}}(a,b,c,d,e,f,g) = \underline{\pi},$$
$$\overline{\mathbf{g}}(\widehat{a},\widehat{b},\widehat{c},\widehat{d},\widehat{e},\widehat{f},\widehat{g}) = \widehat{\pi} \xrightarrow[J \to +\infty]{a.s.} \overline{\mathbf{g}}(a,b,c,d,e,f,g) = \overline{\pi}.$$

Therefore, our simple method of moments estimators provide a consistent estimated identified set of the identification interval $[\pi, \overline{\pi}]$ of polarization index π .

3.B.5 Proof of Proposition 3.2

The proof consists in applying (i) the multivariate Central Limit Theorem to the relevant vector of empirical means, that is the vector of size 7, $(\hat{a}, \hat{b}, \hat{c}, \hat{d}, \hat{e}, \hat{f}, \hat{g})^{\top}$;²² (ii) the delta method with the map $\mathbf{g} := (\mathbf{g}, \overline{\mathbf{g}})^{\top}$ defined into \mathbb{R}^2 .

To lighten notations only, we define the following real random variable

$$L := \frac{K^R K^D}{K - 1} \mathbb{1}\{K > 1\}.$$

To clarify the proper definition of L, remember the spirit of our various conventions: null indicators prevails over a related null denominator, related meaning that the denominator is null if and only if the indicator is switched off. Here, when K = 1, L is thus set to 0.

Central Limit Theorem To apply the multivariate CLT to $(\hat{a}, \hat{b}, \hat{c}, \hat{d}, \hat{e}, \hat{f}, \hat{g})$, we need two conditions: i.i.d. sampling, and a finite variance-covariance matrix.

As explained in Appendix 3.B.4, the non-negative real random variables K^R , K^D , L, $K^R \mathbb{1}\{K=1\}$, and $K^D \mathbb{1}\{K=1\}$ are upper bounded by the variable K. Therefore, the condition $\mathbb{E}[K^2] < +\infty$ implies that those variables also admit a finite second moment. Moreover, the variable $\mathbb{1}\{K=1\}$ has bounded support, hence has a finite variance too. Therefore, under the assumptions of Proposition 3.2, we have the existence of the relevant finite variance-covariance matrix:

$$\Sigma := \mathbb{V}\Big[(K, K^R, K^D, L, K^R \mathbb{1}\{K = 1\}, K^D \mathbb{1}\{K = 1\}, \mathbb{1}\{K = 1\})^\top \Big].$$

Combined with i.i.d. sampling from Assumption 3.1, the multivariate CLT gives

$$\sqrt{J}\left[(\hat{a}, \hat{b}, \hat{c}, \hat{d}, \hat{e}, \hat{f}, \hat{g})^{\top} - (a, b, c, d, e, f, g)^{\top} \right] \xrightarrow[J \to +\infty]{d} \mathcal{N}(0, \Sigma).$$
(3.39)

Delta method We use the delta method with the map $g := (g, \overline{g})$ for

$$g(\hat{a},\hat{b},\hat{c},\hat{d},\hat{e},\hat{f},\hat{g}) = (\underline{\hat{\pi}},\overline{\hat{\pi}})^{\top}, \ g(a,b,c,d,e,f,g) = (\underline{\pi},\overline{\pi})^{\top}.$$

Under our convention for fractions with possible null denominator provided the numerator is null too, and with Assumption 3.1 and $q \in (0, 1)$ that imply $\mathbb{E}[K^R]$ and $\mathbb{E}[K^D]$ are positive, there is no issue of definition of g. In addition, the map is differentiable and we denote by $\mathbf{Jg}|_{(a,b,c,d,e,f,g)}$ the 2 × 7 Jacobian matrix of g evaluated at the vector (a, b, c, d, e, f, g). Therefore, the delta method applied to $g(\cdot)$ and Equation (3.39) gives

$$\sqrt{J}\left[\left(\begin{array}{c} \frac{\widehat{\pi}}{\widehat{\pi}} \end{array}\right) - \left(\begin{array}{c} \frac{\pi}{\pi} \end{array}\right)\right] \xrightarrow[J \to +\infty]{d} \mathcal{N}\left(0, \left(\begin{array}{c} \omega & \tau \\ \tau & \overline{\omega} \end{array}\right)\right),$$

with

$$\begin{pmatrix} \underline{\omega} & \tau \\ \tau & \overline{\omega} \end{pmatrix} := \mathbf{J} \mathbf{g}|_{(a,b,c,d,e,f,g)} \Sigma \mathbf{J} \mathbf{g}|_{(a,b,c,d,e,f,g)}^{\top}.$$
(3.40)

²²The symbol \cdot^{\top} denotes the transpose of a vector or matrix.

Expressions of the asymptotic variance For the practical target of building confidence interval (CI) for π , it remains to compute and consistently estimate the entries, $\underline{\omega}$, τ , and $\overline{\omega}$, of the asymptotic variance-covariance matrix. The analytical derivation of g and the matrix computation of Equation (3.40) does not pose any difficulty but are tedious. We perform those thanks to a computer algebra system, Maxima, and report the results below for completeness. We would be happy to share the Maxima code upon request.

The part related to the upper bound is:

$$\overline{\omega} := \frac{\mathbb{E}[K]^2 \mathbb{V}[L]}{4\mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2} + \frac{\mathbb{E}[K]^2 \mathbb{E}[L]^2 \mathbb{V}\left[K^R\right]}{4\mathbb{E}[K^D]^2 \mathbb{E}[K^R]^4} + \frac{\mathbb{E}[K]^2 \mathbb{E}[L]^2 \mathbb{V}\left[K^D\right]}{4\mathbb{E}[K^D]^4 \mathbb{E}[K^R]^2} + \frac{\mathbb{E}[L]^2 \mathbb{V}[K]}{4\mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2} - \frac{\mathbb{Cov}(K, K^D) \mathbb{E}[K] \mathbb{E}[L]^2}{2\mathbb{E}[K^D]^3 \mathbb{E}[K^R]^2} - \frac{\mathbb{Cov}(K, K^R) \mathbb{E}[K] \mathbb{E}[L]^2}{2\mathbb{E}[K^D]^2 \mathbb{E}[K^R]^3} + \frac{\mathbb{Cov}(K, L) \mathbb{E}[K] \mathbb{E}[L]}{2\mathbb{E}[K^D]^3 \mathbb{E}[K^R]^3} - \frac{\mathbb{Cov}(K^R, L) \mathbb{E}[K] \mathbb{E}[L]}{2\mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2} - \frac{\mathbb{Cov}(K^R, L) \mathbb{E}[K] \mathbb{E}[L]}{2\mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2} - \frac{\mathbb{Cov}(K^R, L) \mathbb{E}[K]^2 \mathbb{E}[L]}{2\mathbb{E}[K^D]^2 \mathbb{E}[K^R]^3} .$$

$$(3.41)$$

The covariance term and the entry related to the lower bound are (far) longer. We express them below as the difference with $\overline{\omega}$. (the expressions are slightly nicer and it is easier to see what happens in the case of point-identification, when $\mathbb{P}(K = 1) = 0$). Other expressions for τ and $\overline{\omega}$ are available in the Maxima script. The asymptotic covariance term τ is defined by

$$\begin{split} &\tau := \overline{\omega} + \\ & \frac{\mathbb{E}[K]^2 \mathbb{E}\left[K^D \mathbb{1}\{K=1\}\right] \mathbb{E}[K^R \mathbb{1}\{K=1\}] \mathbb{E}[L] \mathbb{V}\left[K^R\right]}{4\mathbb{E}[\mathbb{1}\{K=1\}] \mathbb{E}[K^D \mathbb{1}\{K=1\}] \mathbb{E}[K^R \mathbb{1}\{K=1\}] \mathbb{E}[L] \mathbb{V}[K^D\right]} + \frac{\mathbb{E}[K]^2 \mathbb{E}\left[K^D \mathbb{1}\{K=1\}\right] \mathbb{E}[K^R \mathbb{1}\{K=1\}] \mathbb{E}[L] \mathbb{V}[K^D\right]}{4\mathbb{E}[\mathbb{1}\{K=1\}] \mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2} + \\ & \frac{\mathbb{E}[K^D \mathbb{1}\{K=1\}] \mathbb{E}[K^R \mathbb{1}\{K=1\}] \mathbb{E}[L] \mathbb{V}[K]}{4\mathbb{E}[\mathbb{1}\{K=1\}] \mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2} - \\ & \frac{\mathbb{E}[K]^2 \mathbb{E}\left[K^D \mathbb{1}\{K=1\}\right] \mathbb{E}[K^R \mathbb{1}\{K=1\}] \mathbb{E}[L] \mathbb{V}[K]}{2\mathbb{E}[\mathbb{1}\{K=1\}] \mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2} - \\ & \frac{\mathbb{E}[K]^2 \mathbb{E}[K^D \mathbb{1}\{K=1\}] \mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2}{2\mathbb{E}[\mathbb{1}\{K=1\}] \mathbb{E}[K^D]^2 \mathbb{E}[K^R]^2} - \\ & \frac{\mathbb{E}[K]^2 \mathbb{E}\left[K^D \mathbb{1}\{K=1\}\right] \mathbb{E}[K^R \mathbb{1}\{K=1\}] \mathbb{E}[K^R \mathbb{1}\{K=1\}] \mathbb{E}[K^D \mathbb{1}\{K=1\}] \mathbb{E}[K^R \mathbb{1}\{K=1]] \mathbb{E}$$

 $\omega := \overline{\omega} +$ $\frac{\mathbb{E}[K]^2 \,\mathbb{E}\left[K^D \mathbbm{1}\{K=1\}\right]^2 \mathbb{V}\left[K^R \mathbbm{1}\{K=1\}\right]}{4\mathbb{E}[\mathbbm{1}\{K=1\}]^2 \,\mathbb{E}\left[K^D\right]^2 \mathbb{E}\left[K^R\right]^2} + \frac{\mathbb{E}[K]^2 \,\mathbb{E}\left[K^D \mathbbm{1}\{K=1\}\right] \mathbb{E}\left[K^R \mathbbm{1}\{K=1\}\right] \mathbb{E}[L] \,\mathbb{V}\left[K^R\right]}{2\mathbb{E}[\mathbbm{1}\{K=1\}] \,\mathbb{E}\left[K^D\right]^2 \mathbb{E}\left[K^R\right]^4} + \frac{\mathbb{E}[K]^2 \,\mathbb{E}\left[K^R \mathbbm{1}\{K=1\}\right] \mathbb{E}\left[K^R \mathbbm{1}$ $\frac{\mathbb{E}[K]^2 \mathbb{E} \left[K^D \mathbf{1} \{K=1\} \right]^2 \mathbb{E} \left[K^R \mathbf{1} \{K=1\} \right]^2 \mathbb{V} \left[K^R \right]}{4 \mathbb{E}[\mathbf{1} \{K=1\}]^2 \mathbb{E} \left[K^D \right]^2 \mathbb{E} \left[K^R \right]^4} + \frac{\mathbb{E}[K]^2 \mathbb{E} \left[K^R \mathbf{1} \{K=1\} \right]^2 \mathbb{V} \left[K^D \mathbf{1} \{K=1\} \right]}{4 \mathbb{E}[\mathbf{1} \{K=1\}]^2 \mathbb{E} \left[K^D \right]^2 \mathbb{E} \left[K^R \right]^2} + \frac{\mathbb{E}[K]^2 \mathbb{E} \left[K^R \mathbf{1} \{K=1\} \right]^2 \mathbb{E} \left[K^R \mathbf{1} \{K=1\} \right]^2 \mathbb{E} \left[K^R \mathbf{1} \{K=1\} \right]}{4 \mathbb{E}[\mathbf{1} \{K=1\}]^2 \mathbb{E} \left[K^R \right]^2} \mathbb{E} \left[K^R \mathbf{1} \{K=1\} \right]^2 \mathbb{E} \left[K^R \mathbf{1} \{K$ $\frac{\mathbb{E}[K]^2 \mathbb{E}\left[K^D \mathbbm{1}\{K=1\}\right] \mathbb{E}\left[K^R \mathbbm{1}\{K=1\}\right] \mathbb{E}[L] \mathbb{V}\left[K^D\right]}{2\mathbb{E}[\mathbbm{1}\{K=1\}] \mathbb{E}\left[K^D\right]^4 \mathbb{E}\left[K^R\right]^2} + \frac{\mathbb{E}[K]^2 \mathbb{E}\left[K^D \mathbbm{1}\{K=1\}\right]^2 \mathbb{E}\left[K^R \mathbbm{1}\{K=1\}\right]^2 \mathbb{V}\left[K^D\right]}{4\mathbb{E}[\mathbbm{1}\{K=1\}]^2 \mathbb{E}\left[K^D\right]^4 \mathbb{E}\left[K^R\right]^2} + \frac{\mathbb{E}[K]^2 \mathbb{E}\left[K^R \mathbbm{1}\{K=1\}\right]^2 \mathbb{E$ $\frac{\mathbb{E}\Big[K^D \mathbbm{1}\{K=1\}\Big] \mathbb{E}\Big[K^R \mathbbm{1}\{K=1\}\Big] \mathbb{E}[L] \mathbb{V}[K]}{2\mathbb{E}\big[\mathbbm{1}\{K=1\}\big]^2 \mathbb{E}\Big[K^D\Big]^2 \mathbb{E}\Big[K^R\Big]^2} + \frac{\mathbb{E}\Big[K^D \mathbbm{1}\{K=1\}\Big]^2 \mathbb{E}\Big[K^R \mathbbm{1}\{K=1\}\Big]^2 \mathbb{V}[K]}{4\mathbb{E}\big[\mathbbm{1}\{K=1\}\big]^2 \mathbb{E}\Big[K^D\Big]^2 \mathbb{E}\Big[K^R\Big]^2}$ $\frac{\mathbb{E}[K]^2 \,\mathbb{E}\left[K^D \mathbf{1}\{K=1\}\right]^2 \,\mathbb{E}\left[K^R \mathbf{1}\{K=1\}\right]^2 \,\mathbb{V}[\mathbf{1}\{K=1\}]}{4 \mathbb{E}[\mathbf{1}\{K=1\}]^4 \,\mathbb{E}\left[K^D\right]^2 \,\mathbb{E}\left[K^R\right]^2} + \frac{\mathbb{E}[K]^2 \,\mathbb{E}\left[K^D \mathbf{1}\{K=1\}\right] \,\mathbb{E}\left[K^R \mathbf{1}\{K=1\}\right] \,\mathbb{E}\left[K^R \mathbf{1}\{K=1\}\right] \,\mathbb{E}[L]}{2 \mathbb{E}[\mathbf{1}\{K=1\}] \,\mathbb{E}\left[K^D\right]^2 \,\mathbb{E}\left[K^R\right]^2} - \frac{1}{2 \mathbb{E}[\mathbf{1}\{K=1\}] \,\mathbb{E}\left[K^R \mathbf{1}\{K=1\}\right] \,\mathbb{E}\left[K^R \mathbf{1}\{K=1\}\right$ $\frac{\mathbb{E}[K]^2 \mathbb{E}\left[K^D \mathbb{1}\{K=1\}\right] \mathbb{E}\left[K^R \mathbb{1}\{K=1\}\right] \mathbb{E}[L]}{\mathbb{E}[\mathbb{1}\{K=1\}] \mathbb{E}\left[K^D\right]^2 \mathbb{E}\left[K^R\right]^2} - \frac{\mathbb{Cov}(K, \mathbb{1}\{K=1\}) \mathbb{E}[K] \mathbb{E}\left[K^D \mathbb{1}\{K=1\}\right] \mathbb{E}\left[K^R \mathbb{1}\{K=1\}\right] \mathbb{E}[L]}{2\mathbb{E}[\mathbb{1}\{K=1\}]^2 \mathbb{E}\left[K^D\right]^2 \mathbb{E}\left[K^R\right]^2} + \frac{\mathbb{E}[\mathbb{E}[K] \mathbb{E}[K] \mathbb{E}[K]$ $\frac{\mathbb{C}\mathrm{ov}\left(\boldsymbol{K}^{D}, \mathbf{1}\{\boldsymbol{K}=1\}\right)\mathbb{E}[\boldsymbol{K}]^{2}\mathbb{E}\left[\boldsymbol{K}^{D}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}[\boldsymbol{L}]}{2\mathbb{E}[\mathbf{1}\{\boldsymbol{K}=1\}]^{2}\mathbb{E}\left[\boldsymbol{K}^{D}\right]^{3}\mathbb{E}\left[\boldsymbol{K}^{R}\right]^{2}} - \frac{\mathbb{C}\mathrm{ov}\left(\boldsymbol{K}, \boldsymbol{K}^{D}\right)\mathbb{E}[\boldsymbol{K}]\mathbb{E}\left[\boldsymbol{K}^{D}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}[\boldsymbol{L}]}{\mathbb{E}[\mathbf{1}\{\boldsymbol{K}=1\}]\mathbb{E}\left[\boldsymbol{K}^{D}\right]^{3}\mathbb{E}\left[\boldsymbol{K}^{R}\right]^{2}} + \frac{\mathbb{E}[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K}^{R}\mathbf{1}\{\boldsymbol{K}=1\}\right]\mathbb{E}\left[\boldsymbol{K$ $\frac{\mathbb{C}\operatorname{ov}\left(K, K^{D}\mathbb{1}\{K=1\}\right)\mathbb{E}[K]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}[L]}{2\mathbb{E}[\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{C}\operatorname{ov}\left(K^{D}, K^{D}\mathbb{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}[L]}{2\mathbb{E}[\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{2}} + \frac{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}$ $\frac{\operatorname{Cov}\left(K^{R}, 1\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}1\{K=1\}\right]\mathbb{E}\left[K^{R}1\{K=1\}\right]\mathbb{E}[L]}{2\mathbb{E}[1\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{3}} - \frac{\operatorname{Cov}\left(K, K^{R}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}1\{K=1\}\right]\mathbb{E}\left[K^{R}1\{K=1\}\right]\mathbb{E}[L]}{\mathbb{E}[1\{K=1\}]\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{3}} - \frac{\operatorname{Cov}\left(K^{R}, K^{D}1\{K=1\}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{E}\left[K^{R}1\{K=1\}\right]\mathbb{E}[L]\right]}{2\mathbb{E}[1\{K=1\}]\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{3}} - \frac{\operatorname{Cov}\left(K^{R}, K^{D}1\{K=1\}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{E}\left[K^{R}1\{K=1\}\right]\mathbb{E}[L]\right]}{2\mathbb{E}[1\{K=1\}]\mathbb{E}\left[K^{D}\mathbb{E}\left[K^{R}1\{K=1\}\right]\mathbb{E}[L]\right]} + \frac{\operatorname{Cov}\left(K^{R}, K^{D}1\{K=1\}\right)\mathbb{E}\left[K^{D}\mathbb{E}\left[K^{R}1\{K=1\}\right]\mathbb{E}[L]\right]}{2\mathbb{E}[1\{K=1\}]\mathbb{E}\left[K^{D}\mathbb{E}\left[K^{R}\mathbb{E}\left$ $\frac{\mathbb{C}\operatorname{ov}\left(K, K^{R}\mathbb{1}\{K=1\}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]\mathbb{E}[L]}{2\mathbb{E}[\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{C}\operatorname{ov}\left(K^{D}, K^{R}\mathbb{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]\mathbb{E}[L]}{2\mathbb{E}[\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}$ $\frac{\mathbb{C}\mathrm{ov}\left(K^{R},K^{R}\mathbbm{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbbm{1}\{K=1\}\right]\mathbb{E}[L]}{2\mathbb{E}[\mathbbm{1}\{K=1\}]\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{3}}-\frac{\mathbb{C}\mathrm{ov}(K,\mathbbm{1}\{K=1\})\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbbm{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbbm{1}\{K=1\}]^{3}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{2}}+$ $\frac{\mathbb{C}\operatorname{ov}\left(K^{D}, \mathbb{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{3}\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{C}\operatorname{ov}\left(K, K^{D}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{C}\operatorname{ov}\left(K, K^{D}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}}{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}}{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}}{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}}{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}}{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{$ $\frac{\mathbb{C}\mathrm{ov}\left(K^{D}\mathbb{1}\{K=1\},\mathbb{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{3}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{2}} + \frac{\mathbb{C}\mathrm{ov}\left(K,K^{D}\mathbb{1}\{K=1\}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]} + \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]}{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]} + \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]}{+ \frac{\mathbb{E}[K^{R}\mathbb{1}\{K$ $\frac{\mathbb{C}\mathrm{ov}\left(K^{D}, K^{D}\mathbbm{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbbm{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{2}} + \frac{\mathbb{C}\mathrm{ov}\left(K^{R}, \mathbbm{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbbm{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbbm{1}\{K=1\}]^{3}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{3}} - \frac{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{\mathbb{E}[\mathbbm{1}\{K=1\}]^{3}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}} - \frac{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}} + \frac{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}} + \frac{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}\left[K^{R}\mathbbm{1}\{K=1\}\right]^{2}}{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}}{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}}{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}}{\mathbb{E}[\mathbbm{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[\mathbbm{1}\{\mathbbm{1}\{K=1\}]^{2}}{\mathbb{E}[\mathbbm{1}\{\mathbbm{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[\mathbbm{1}\{\mathbbm{1}\{K=1\}]^{2}}{\mathbb{E}[\mathbbm{1}\{\mathbbm{1}\{K=1\}]^{2}} + \frac{\mathbb{E}[\mathbbm{1}\{\mathbbm{1}\{\mathbbm{1}\{K=1\}\}}{+ \frac{\mathbb{E}[\mathbbm{1}\{\mathbbm{1}\{\mathbbm{1}\{\mathbbm{1}\{\mathbbm{1}\{K=1\}\}}{+ \frac{\mathbb{E}[\mathbbm{1}\{\mathbbm{1$ $\frac{\mathbb{C}\mathrm{ov}\left(K,K^{R}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{3}} + \frac{\mathbb{C}\mathrm{ov}\left(K^{R},K^{D}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]^{2}}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{3}}$ $\frac{\mathbb{C}\mathrm{ov}\Big(K^{R}, K^{D}\mathbb{1}\{K=1\}\Big)\mathbb{E}[K]^{2}\mathbb{E}\Big[K^{D}\mathbb{1}\{K=1\}\Big]\mathbb{E}\Big[K^{R}\mathbb{1}\{K=1\}\Big]^{2}}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\Big[K^{D}\Big]^{2}\mathbb{E}\Big[K^{R}\Big]^{3}} - \frac{\mathbb{C}\mathrm{ov}\Big(K^{R}\mathbb{1}\{K=1\},\mathbb{1}\{K=1\}\Big)\mathbb{E}[K^{D}\mathbb{1}^{2}\mathbb{E}\Big[K^{D}\mathbb{1}^{2}\mathbb{E}\Big[K^{R}\mathbb{1}^{2}\mathbb{1}^{2}\mathbb{E}\Big[K^{R}\mathbb{1}^{2}\mathbb{1}^{2}\mathbb{E}\Big[K^{R}\mathbb{1}^{2}\mathbb{1}^{2}\mathbb{E$ $\frac{\mathbb{C}\operatorname{ov}\left(K, K^{R}\mathbb{1}\{K=1\}\right)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{C}\operatorname{ov}\left(K^{D}, K^{R}\mathbb{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{2}} + \frac{\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]^{2}}$ $\frac{\mathbb{C}\mathrm{ov}\left(K^{R}\mathbb{1}\{K=1\},K^{D}\mathbb{1}\{K=1\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{2\mathbb{E}[\mathbb{1}\{K=1\}]^{2}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{2}} + \frac{\mathbb{C}\mathrm{ov}(K,L)\mathbb{E}[K]\mathbb{E}\left[K^{D}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{2\mathbb{E}[\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}]}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}\right]}{\mathbb{E}\left[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[K^{R}\mathbb{1}\{K=1\}\right]}{- \frac{E}[K^{R}\mathbb{1}\{K=1\}\right]} - \frac{\mathbb{E}[$ $\frac{\mathbb{C}\mathrm{ov}\left(K^{D},L\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\left\{K=1\right\}\right]\mathbb{E}\left[K^{R}\mathbb{1}\left\{K=1\right\}\right]}{2\mathbb{E}[\mathbb{1}\left\{K=1\right\}]\mathbb{E}\left[K^{D}\right]^{3}\mathbb{E}\left[K^{R}\right]^{2}} - \frac{\mathbb{C}\mathrm{ov}\left(K^{R},K^{R}\mathbb{1}\left\{K=1\right\}\right)\mathbb{E}[K]^{2}\mathbb{E}\left[K^{D}\mathbb{1}\left\{K=1\right\}\right]^{2}\mathbb{E}\left[K^{R}\mathbb{1}\left\{K=1\right\}\right]}{2\mathbb{E}[\mathbb{1}\left\{K=1\right\}]^{2}\mathbb{E}\left[K^{D}\right]^{2}\mathbb{E}\left[K^{R}\right]^{3}}$ $\frac{\mathbb{C}\mathrm{ov}\left(K^{R},L\right)\mathbb{E}[K]^{2}\,\mathbb{E}\Big[K^{D}\mathbbm{1}\{K=1\}\Big]\,\mathbb{E}\Big[K^{R}\mathbbm{1}\{K=1\}\Big]}{2\mathbb{E}[\mathbbm{1}\{K=1\}]\,\mathbb{E}\Big[K^{D}\Big]^{2}\,\mathbb{E}\Big[K^{R}\Big]^{3}}\;.$ (3.43)

The asymptotic variance for the lower bound $\underline{\omega}$ is defined by

Those expressions call for several comments. First, as stated in Proposition 3.2 and despite their length, the three quantities $\underline{\omega}$, τ , and $\overline{\omega}$ are indeed continuous functions of the expectations and variances of, or covariances between, the concerned random variables, namely K, K^R , K^D , $K^R K^D \mathbb{1}\{K = 1\}/(K - 1)$, $K^R \mathbb{1}\{K = 1\}$, $K^D \mathbb{1}\{K = 1\}$, and $\mathbb{1}\{K = 1\}$. Moreover, those variables are observed in the data, so that simple and consistent estimators by the method of moments are available and can be used to construct confidence intervals.²³

Second, two examinations ensure the proper definition of the asymptotic variancecovariance matrix. In $\overline{\omega}$, the expectations $\mathbb{E}[K^R]$ and $\mathbb{E}[K^D]$ constitute all the denominators. Under Assumption 3.1 and $q \in (0, 1)$, the expectations are positive, hence the proper definition of $\overline{\omega}$. Furthermore, the quantities $\tau - \overline{\omega}$ and $\underline{\omega} - \overline{\omega}$ share a noteworthy feature: they are entirely made up of fractions whose common characteristic is to contain either $\mathbb{E}[K^R \mathbb{1}\{K=1\}]/\mathbb{E}[\mathbb{1}\{K=1\}]$ or $\mathbb{E}[K^D \mathbb{1}\{K=1\}]/\mathbb{E}[\mathbb{1}\{K=1\}]$. That consideration establishes the proper definition of τ and $\overline{\omega}$. Indeed, their denominators are made up of $\mathbb{E}[\mathbb{1}\{K=1\}], \mathbb{E}[K^R]$, and $\mathbb{E}[K^D]$. If $\mathbb{E}[\mathbb{1}\{K=1\}] = \mathbb{P}(K=1) > 0$, combined with $\mathbb{E}[K^R]$ and $\mathbb{E}[K^D]$ positive, the denominators are not null. On the contrary, if there is no single-choice options, $\mathbb{P}(K=1) = 0$, our convention of Theorem 3.2 applies and, as logically expected, those fractions are null and, consequently, the differences vanish: $\overline{\omega} = \tau = \underline{\omega}$. In that point-identification setting, the two bounds, $\underline{\pi} = \overline{\pi} = \pi$, and the two estimators, $\hat{\underline{\pi}} = \hat{\overline{\pi}}$, coincide. Then, Proposition 3.2 reduces to

$$\sqrt{J}\left(\widehat{\pi}-\pi\right) \xrightarrow[J \to +\infty]{d} \mathcal{N}\left(0,\overline{\omega}\right).$$
 (3.44)

Thus, if $\mathbb{P}(K = 1) = 0$, our method yields point-identification and the classical result of a consistent, asymptotically normal, point-estimator of the parameter of interest π .

Third, alternative, shorter, linearized expressions of the asymptotic variance-covariance matrix of $(\hat{\pi}, \hat{\pi})$ are $\underline{\omega} = \mathbb{V}(\underline{\delta})$, $\tau = \mathbb{C}\operatorname{ov}(\underline{\delta}, \overline{\delta})$, $\overline{\omega} = \mathbb{V}(\overline{\delta})$, where the real random variables δ and $\overline{\delta}$ are defined by:

$$\begin{split} \overline{\delta} &:= (\overline{\pi} - 1) \left\{ \frac{K}{\mathbb{E}[K]} + \frac{K^R K^D \mathbbm{1}\{K > 1\} / (K - 1)}{\mathbb{E}[K^R K^D \mathbbm{1}\{K > 1\} / (K - 1)]} - \frac{K^R}{\mathbb{E}[K^R]} - \frac{K^D}{\mathbb{E}[K^D]} \right\}, \\ \underline{\delta} &:= \overline{\delta} - \frac{\mathbb{E}[K] \mathbb{E}\Big[K^R \mathbbm{1}\{K = 1\}\Big] \mathbb{E}\Big[K^D \mathbbm{1}\{K = 1\}\Big]}{2\mathbb{E}[K^R] \mathbb{E}[K^D] \mathbb{E}[\mathbbm{1}\{K = 1\}]} \left\{ \frac{K}{\mathbb{E}[K]} + \frac{K^R \mathbbm{1}\{K = 1\}}{\mathbb{E}[K^R \mathbbm{1}\{K = 1\}]} + \frac{K^D \mathbbm{1}\{K = 1\}}{\mathbb{E}[K^D \mathbbm{1}\{K = 1\}]} - \frac{K^R}{\mathbb{E}[K^R]} - \frac{K^D}{\mathbb{E}[K^D]} - \frac{\mathbbm{1}\{K = 1\}}{\mathbb{E}[\mathbbm{1}\{K = 1\}]} \right\}. \end{split}$$

The point-identification setting is more visible in those expressions. If $\mathbb{P}(K=1) = 0$, the convention of Theorem 3.2 entails that the first fraction in $\underline{\delta}$ is null, so that $\underline{\delta} = \overline{\delta}$ and $\underline{\omega} = \tau = \overline{\omega} = \mathbb{V}(\overline{\delta})$.

²³The expressions appear long and ugly but provide straightforward method of moments estimators, virtually costless to compute. In the end, the formal expressions are not very important; we compute them in Maxima and then copy-and-paste them twice before forgetting them. One copy is in the .tex that generates this documents. The second appears in the R code to compute the consistent estimators. In the code, we use simplifications coming from $K := K^R + K^D$ to save computational time. For instance, to obtain the empirical counterparts of $\mathbb{E}[K]$, $\mathbb{E}[K^R]$, and $\mathbb{E}[K^D]$, there is no need to compute three sample means over, possibly, millions of options; two are enough, plus an immediate difference between those two real numbers.

3.B.6 Proof of Proposition 3.3

The proof consists in checking Assumption 1(i) and (ii) in Stoye (2009). Indeed, by construction of our bound estimators, $\{\hat{\pi} \geq \hat{\pi}\}$ holds almost surely for any DGP satisfying Assumption 3.1. Therefore, Lemma 3 of Stoye (2009) gives Assumption 3 with the notation of this paper, then Proposition 1 gives the result.

The first part of 1(ii) (the asymptotic variance of the lower and upper bound estimators are bounded by finite and positive constants uniformly over possible DGP of the statistical model) is assumed in Proposition 3.3.²⁴

The second part of 1(ii) states that the identification interval's length is bounded uniformly over possible DGP. This is indeed the case from Equation 3.3, that we recall here

$$\Delta := \overline{\pi} - \underline{\pi} = \mathbb{E} \Big[K^R \, | \, K = 1 \Big] \, \mathbb{E} \Big[K^D \, | \, K = 1 \Big] \, \mathbb{P}(K = 1) \, .$$

Since $K^R \leq K$, $K^D \leq K$ and a probability cannot exceed 1, we have $\Delta \leq 3$ for any DGP under Assumption 3.1.

Assumption 1(i) requires joint asymptotic normality of the bound estimators $(\hat{\pi}, \hat{\pi})$ uniformly over the possible DGP as well as estimators $\hat{\omega}, \hat{\tau}$ and $\hat{\omega}$ of the asymptotic variance-covariance matrix that are uniformly consistent. We already have those results pointwise (Proposition 3.2 for the asymptotic normality and consistent plug-in estimators of the asymptotic variance-covariance of $(\hat{\pi}, \hat{\pi})$; see details in the next paragraph). The extension to uniformly valid results can be obtained from the results of Romano (2004) (R) and Kasy (2019) (K). For the uniform asymptotic normality, Lemma 1 of R gives a uniform Central Limit Theorem (CLT) on all needed sample means (see the quantities a to g in Section 3.B.4; remember that all the relevant variables are upper bounded by K, so the technical condition on the control of its $2 + \varepsilon$ moment is enough). Then the use of the uniform delta method (Theorem 2) of K yields a uniform CLT for the bound estimators. For the uniform consistency of the estimators of $\hat{\omega}, \bar{\omega}$ and τ , the proof is symmetric: Lemma 2 in R gives uniform convergence in probability on all required sample means, and the uniform continuous mapping theorem (Theorem 1) of K enables to deduce a uniform consistency of the asymptotic variance-covariance estimators.

Consistent estimators for the asymptotic variance The inspection of Equations (3.43), (3.42), and (3.41) gives natural estimators through the method of moments. Indeed, $\underline{\omega}$, τ , and $\overline{\omega}$ write as continuous functions of expectations, covariances, and variances of observed variables. We introduce the following notations to express the replacement of population moments by empirical counterparts: for any random variable V and W among K, K^R , K^D , $L := K^R K^D \mathbb{1}\{K = 1\}/(K - 1), K^R \mathbb{1}\{K = 1\},$ $K^D \mathbb{1}\{K = 1\}$, and $\mathbb{1}\{K = 1\}$,

$$\widehat{\mathbb{E}}[V] := \frac{1}{J} \sum_{j=1}^{J} V_j, \quad \widehat{\mathbb{V}}[V] := \frac{1}{J} \sum_{j=1}^{J} \left(V_j - \widehat{\mathbb{E}}[V] \right)^2,$$

$$\widehat{\mathbb{Cov}}(V, W) := \frac{1}{J} \sum_{j=1}^{J} \left(V_j - \widehat{\mathbb{E}}[V] \right) \left(W_j - \widehat{\mathbb{E}}[W] \right). \quad (3.45)$$

²⁴Within our DGP, provided adequate moment conditions on K, it should probably be possible to obtain the result without further assumptions. We leave this task for future work.

The estimators $\underline{\hat{\omega}}$, $\hat{\tau}$ and $\overline{\hat{\omega}}$ are defined respectively by Equations (3.43), (3.42), and (3.41) when replacing $\mathbb{E}[\cdot]$, $\mathbb{V}[\cdot]$, and $\mathbb{C}ov(\cdot, \cdot)$ by their empirical counterparts, $\widehat{\mathbb{E}}[\cdot]$, $\widehat{\mathbb{V}}[\cdot]$, and $\widehat{\mathbb{C}ov}(\cdot, \cdot)$.

Intuition on the form of our CI The joint asymptotic normality of the estimators of the bounds and consistent estimators of their asymptotic variance-covariance matrix constitute the two ingredients for our CI on π defined by Equations (3.6) and (3.7). The specific form of $\text{CI}_{1-\alpha}^{\pi}$ comes from adaptation to the partial identification of the polarization index. Following the intuitions of Imbens and Manski (2004) and Stoye (2009), for partially identified parameters, we expect that, asymptotically, the length of the identification interval, here $\Delta := \overline{\pi} - \underline{\pi}$, will be large relative to sampling uncertainty. Crudely put, π will be either above the upper bound or below the lower bound. If we accept non-coverage to occur with an asymptotic probability equal to $\alpha \in (0, 1)$, that idea suggests using the quantile $q_{\mathcal{N}(0,1)}(1-\alpha)$, instead of the usual $q_{\mathcal{N}(0,1)}(1-\alpha/2)$. Indeed, the non-coverage risk is one-sided only. Nevertheless, the previous intuition fails when Δ does not diverge relative to sampling uncertainly, that is, when the latter does not become negligible compared to the uncertainty coming from partial identification. The trick is to choose a critical value that balances sampling and partial identification uncertainties, hence the definition of $q(\alpha)$ that we recall and justify below.

Let $\widehat{\Delta} := \widehat{\pi} - \widehat{\pi}$ and $\Phi(\cdot)$ denote the cumulative distribution function of the standard Gaussian $\mathcal{N}(0,1)$ distribution. For $\alpha \in (0,1)$, we define $q(\alpha)$ as the solution of

$$\Phi\left(\mathbf{q}(\alpha) + \frac{\sqrt{J}\widehat{\Delta}}{\max\left\{\sqrt{\widehat{\omega}}, \sqrt{\widehat{\omega}}\right\}}\right) - \Phi(-\mathbf{q}(\alpha)) = 1 - \alpha, \tag{3.46}$$

Then, for the reminder, our CI writes

$$\operatorname{CI}_{1-\alpha}^{\pi} := \left[\underline{\widehat{\pi}} - \frac{q(\alpha)\sqrt{\widehat{\omega}}}{\sqrt{J}}, \overline{\widehat{\pi}} + \frac{q(\alpha)\sqrt{\widehat{\omega}}}{\sqrt{J}} \right]$$

The equation (3.46) does define a unique $q(\alpha) \in (0, +\infty)$. Indeed, for a given sample and corresponding estimates, $\sqrt{J}\widehat{\Delta}/\max\left\{\sqrt{\widehat{\omega}}, \sqrt{\widehat{\omega}}\right\}$ is a constant, and the function $z \in \mathbb{R} \mapsto \Phi(z + \text{constant}) - \Phi(-z)$ has the following properties: continuous, strictly increasing, converges to 1 when $z \to +\infty$, and non-positive when evaluated at 0 since $\Phi(\cdot)$ is maximal at this point. Therefore, as $1 - \alpha \in (0, 1)$, there is a unique solution $q(\alpha)$ to the equation.

The resulting critical value balances sampling uncertainty and the one stemming from partial identification. In case of point-identification, with $\hat{\Delta} = 0$, remark that the aforementioned equation becomes

$$\Phi(\mathbf{q}(\alpha)) - \Phi(-\mathbf{q}(\alpha)) = 1 - \alpha \iff \mathbf{q}(\alpha) = \Phi^{-1}(1 - \alpha/2) =: \mathbf{q}_{\mathcal{N}(0,1)}(1 - \alpha/2),$$

by symmetry of the distribution: for any real number z, $\Phi(-z) = 1 - \Phi(z)$.

Thus, as already described in Equation (3.44), in cases of point-identification, $\operatorname{CI}_{1-\alpha}^{\pi}$ boils down to the usual symmetric CI obtained with an asymptotically normal point-estimator: $\left[\widehat{\pi} \pm q_{\mathcal{N}(0,1)}(1-\alpha/2)\sqrt{\widehat{\omega}}/\sqrt{J}\right]$. The more sampling uncertainty prevails over the uncertainty of partial identification, namely the larger the denominator $\sqrt{\widehat{\omega}}$ or $\sqrt{\widehat{\omega}}$ relative to $\widehat{\Delta}$, the closer is $\operatorname{CI}_{1-\alpha}^{\pi}$ to that classical CI.

On the contrary, when the bulk of uncertainty lies in partial identification rather than in sampling error, that is, when $\hat{\Delta}$ divided by $J^{-1/2} \max\left\{\sqrt{\hat{\omega}}, \sqrt{\hat{\omega}}\right\}$ diverges to $+\infty$, the critical values solves

$$\Phi(+\infty) - \Phi(-\mathbf{q}(\alpha)) = 1 - \alpha \iff \mathbf{q}(\alpha) = \Phi^{-1}(1 - \alpha) =: \mathbf{q}_{\mathcal{N}(0,1)}(1 - \alpha),$$

where the slight abuse of notation $\Phi(+\infty)$ stands for the limit in $+\infty$ of $\Phi(\cdot)$, namely 1. In that situation, the initial intuition of a one-sided risk of non-coverage applies and we can afford a smaller critical value $q_{\mathcal{N}(0,1)}(1-\alpha) < q_{\mathcal{N}(0,1)}(1-\alpha/2)$.

In between those two limit cases, the definition of $q(\alpha)$ in Equation (3.46) trade off the two sources of uncertainty.

3.B.7 Special case of independence between K and ρ

Beyond partial identification, this article proposes two methods that, under additional assumptions, yield point-identification of π . They are notably useful when the proportion of options chosen only once are naturally high in the analyzed data-generating process, so that the identified bounds might not be informative on the level of polarization.

The main body of the article (Section 3.3.2) presents a method based on the extrapolation of m(1), the sole unidentified part of π . This appendix develops an alternative method that requires some independence between the random variables K and ρ , namely $\mathbb{C}\mathrm{ov}(K,\rho) = \mathbb{C}\mathrm{ov}(K,\rho^2) = \mathbb{C}\mathrm{ov}(K^2,\rho^2) = 0$. Under that restriction, π is point-identified, and the method of moments provides a computationally light estimator, consistent and asymptotically normal.

We present those results in appendix because the null covariances assumption might be restrictive in practice. The credibility of that additional assumption depends on the application. In our case, it posits that the popularity of a word, the number of occurrences K, is unrelated with the fact that the word is or not a partian expression, as measured by its conditional probability ρ to be said by a Republican when pronounced – the closer ρ to 0 or 1, the more partian can be deemed the expression. If political speech matters and polarization exists, it is sensible that congresspeople do choose the words they use and that the most partian phrases are more pronounced. If so, the hypothesis of null covariance between K and ρ is violated.

Intuition As discussed after Theorem 3.2, the options chosen only once are the the source of partial identification of π . The idea is that, for K = 1, we are unable to determine whether the option is more or less chosen by one group instead of the other. Formally, it happens that, for any $k \in \mathbb{N}^*$, the binomial assumption $K^R | K, \rho \sim \text{Binomial}(K, \rho)$ gives the identification of the first k moments of the distribution $\mathbb{P}^{\rho | K = k}$.²⁵ However, π involves the second moment of ρ . Hence, when $\mathbb{P}(K = 1) > 0$, we can only identify the first moment of $\mathbb{P}^{\rho | K = 1}$ and it is impossible to reconstruct the second moment of ρ , hence the partial identification of π .

The additional assumption between K and ρ states some independence between the two random variables. It enables to write the index π as a function of the distribution of ρ only, in opposition the joint distribution $P^{(K,\rho)}$. Actually, under that restriction, π is even a function of the first two moments of ρ , $\mathbb{E}[\rho]$ and $\mathbb{E}[\rho^2]$. Then, thanks to the binomial assumption, those moments are identified whenever we observe options chosen more than once. It results in point-identification of π .

²⁵We refer to Equation (2.2), p45 of D'Haultfœuille and Rathelot (2017).

Identification The following theorem formalizes the previous reasoning.

Theorem 3.3 (Point-identification). If Assumption 3.1 holds, $\mathbb{P}(K > 1) > 0$, and $\mathbb{C}ov(K, \rho) = \mathbb{C}ov(K, \rho^2) = \mathbb{C}ov(K^2, \rho^2) = 0$, then

$$\pi = 1 - \frac{\mathbb{E}[\rho] - \mathbb{E}[\rho^2]}{2\mathbb{E}[\rho] \left(1 - \mathbb{E}[\rho]\right)},$$

and

$$\mathbb{E}[\rho] = \frac{\mathbb{E}\left[K^R\right]}{\mathbb{E}[K]}, \quad \mathbb{E}\left[\rho^2\right] = \frac{\mathbb{E}\left[K^R(K^R - 1)\right]}{\mathbb{E}[K(K - 1)]}$$

Consequently, since our DGP identifies $P^{(K^R,K)}$, the polarization index π is point-identified under those assumptions.

Proof of Theorem 3.3. To begin with, the assumption $\mathbb{P}(K > 1) > 0$ is equivalent to $\mathbb{E}[K(K-1)] > 0$ since the variable K(K-1) is non-negative, which ensures the denominator is not null in the expression of $\mathbb{E}[\rho^2]$. It requires there are some options chosen at least twice. It is a mild requirement in practice and, otherwise, it is hopeless to study polarization anyway.

The definition of π gives

$$\pi := 1 - \frac{\mathbb{E}[K\rho(1-\rho)]}{2\mathbb{E}[K]\,q(1-q)} \stackrel{\text{EL}}{=} 1 - \frac{\mathbb{E}[K\rho] - \mathbb{E}\left[K\rho^2\right]}{2\mathbb{E}[K]\,q(1-q)}$$

The idea is then to use the assumed null covariances to simplify the expression and get point-identification.

Regarding q, we have

$$q := \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]} = \frac{\mathbb{E}[K]\mathbb{E}[\rho]}{\mathbb{E}[K]} = \mathbb{E}[\rho],$$

where the second equality comes from $\mathbb{C}ov(K, \rho) = 0$. Besides, $\mathbb{E}[K\rho] = \mathbb{E}[K^R]$. Hence, $q = \mathbb{E}[\rho]$ is identified by $\mathbb{E}[K^R]/\mathbb{E}[K]$.

For the numerator, using $Cov(K, \rho) = Cov(K, \rho^2) = 0$, we have

$$\mathbb{E}[K\rho] - \mathbb{E}[K\rho^2] = \mathbb{E}[K] \mathbb{E}[\rho] - \mathbb{E}[K] \mathbb{E}[\rho^2] = \mathbb{E}[K] \left(\mathbb{E}[\rho] - \mathbb{E}[\rho^2]\right)$$

Combining the last three equations yields the expression of π stated in Theorem 3.3.

It remains to identify the quantity $\mathbb{E}[\rho^2]$. The start is similar to the general case of Theorem 3.2. Using our binomial assumption, we can show

$$\mathbb{E}\left[K^{R}(K^{R}-1) \mid K, \rho\right] = K(K-1)\rho^{2}$$

We then take the expectation over the joint distribution $P^{(K,\rho)}$ of that equality:

$$\mathbb{E}\Big[\mathbb{E}\Big(K^R(K^R-1) \,|\, K, \rho\Big)\Big] = \mathbb{E}\Big[K(K-1)\rho^2\Big]$$

By the Law of Iterated Expectations, the left-hand side is equal to $\mathbb{E}[K^R(K^R-1)]$. For the right-hand side, we have

$$\mathbb{E}\left[K(K-1)\rho^2\right] \stackrel{\text{EL}}{=} \mathbb{E}\left[K^2\rho^2\right] - \mathbb{E}\left[K\rho^2\right]$$
$$= \mathbb{E}\left[K^2\right] \mathbb{E}\left[\rho^2\right] - \mathbb{E}[K] \mathbb{E}\left[\rho^2\right]$$
$$\stackrel{\text{EL}}{=} \mathbb{E}\left[K^2 - K\right] \mathbb{E}\left[\rho^2\right]$$
$$= \mathbb{E}[K(K-1)] \mathbb{E}\left[\rho^2\right],$$

where the second equality uses $\mathbb{C}ov(K, \rho^2) = \mathbb{C}ov(K^2, \rho^2) = 0$. Finally, dividing by $\mathbb{E}[K(K-1)]$ which is positive, we obtain

$$\mathbb{E}\left[\rho^{2}\right] = \frac{\mathbb{E}\left[K^{R}(K^{R}-1)\right]}{\mathbb{E}[K(K-1)]}.$$

This concludes the proof. We see that the null covariances stated in the assumption are sufficient. They are implied by the independence between K and ρ .

Estimation As for the identified bounds, the result of Theorem 3.3 suggests a straightforward estimator of π by the method of moments:

$$\widehat{\pi}_{\text{point}} := 1 - \frac{\frac{\sum_{j=1}^{J} K_{j}^{R}}{\sum_{j=1}^{J} K_{j}} - \frac{\sum_{j=1}^{J} K_{j}^{R} (K_{j}^{R} - 1)}{\sum_{j=1}^{J} K_{j} (K_{j} - 1)}}{2 \frac{\sum_{j=1}^{J} K_{j}^{R}}{\sum_{j=1}^{J} K_{j}} \frac{\sum_{j=1}^{J} K_{j}^{D}}{\sum_{j=1}^{J} K_{j}}}.$$
(3.47)

Replacing expectations by sample means, the J^{-1} cancels out. Moreover, we use $K_j - K_i^R = K_j^D$ in the estimation of $1 - \mathbb{E}[\rho]$.

3.B.8 Naive method and small-unit bias

This appendix presents a competing naive method to identify and estimate π . Instead of relying on our binomial assumption, it replaces the unobserved probability ρ by its empirical counterpart, the proportion K^R/K . We show that it leads to overestimate the real level of polarization, as quantified by the index π .

Naive method From the definitions of q and π in Equation (3.2), using only the linearity of expectation, we have

$$\pi = 1 - \frac{\mathbb{E}[K\rho] - \mathbb{E}[K\rho^2]}{2\mathbb{E}[K] \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]} \left(1 - \frac{\mathbb{E}[K\rho]}{\mathbb{E}[K]}\right)}.$$

The underlying probability ρ is unobserved. However, a natural estimator exists and is observed in the data: K^R/K . What happens if we use that proportion?

We define π_{naive} simply by replacing ρ by K^R/K in the previous equation:

$$\pi_{\text{naive}} := 1 - \frac{\mathbb{E}\left[K^R\right] - \mathbb{E}\left[(K^R)^2/K\right]}{2\mathbb{E}[K] \frac{\mathbb{E}\left[K^R\right]}{\mathbb{E}[K]} \left(1 - \frac{\mathbb{E}\left[K^R\right]}{\mathbb{E}[K]}\right)}.$$
(3.48)

Remark that the expression involves dividing by K, a variable whose support is \mathbb{N} a priori, thus possibly equal to 0. Nonetheless, there is no issue. As stated in Theorem 3.2 and detailed in Appendix 3.B.2, without loss of generality as regards identification and estimation, we can assume $\mathbb{P}(K > 0) = 1$, that is, we can ignore options with zero choices, for which K is null.

A noteworthy fact is that, in doing so, we have

$$\mathbb{E}[K\rho] \xrightarrow{\text{replaced by}} \mathbb{E}\left[K\frac{K^R}{K}\right] = \mathbb{E}\left[K^R\right].$$

In this respect, the naive replacement preserves the equality $\mathbb{E}[K\rho] = \mathbb{E}[K^R]$ formally derived under our statistical model trough the binomial assumption (see Equation (3.16)). In other words, by chance, the naive method succeeds in identifying $\mathbb{E}[K\rho]$.

Upward bias From its definition in Equation 3.2 and using $\mathbb{E}[K\rho] = \mathbb{E}[K^R]$, the parameter π can be written

$$\pi = 1 - \frac{\mathbb{E}\left[K^R\right] - \mathbb{E}\left[K\rho^2\right]}{2\mathbb{E}[K] \frac{\mathbb{E}\left[K^R\right]}{\mathbb{E}[K]} \left(1 - \frac{\mathbb{E}\left[K^R\right]}{\mathbb{E}[K]}\right)}$$

The comparison of π in Equation (3.18) with π_{naive} in Equation (3.48) reveals that the only divergence is the difference between $\mathbb{E}[K\rho^2]$, that appears in the former, and $\mathbb{E}[(K^R)^2/K]$ in the latter. Unlike $\mathbb{E}[K\rho]$, there is no hazardous unfounded equality here. On the contrary, Jensen's inequality implies $\mathbb{E}[(K^R)^2/K] \geq \mathbb{E}[K\rho^2]$. Hence, it shows that $\pi_{\text{naive}} \geq \pi$, that is, the naive method overestimates the real level of polarization.

$$\mathbb{E}\left[K\rho^{2}\right] \stackrel{\text{replaced by}}{\longrightarrow} \mathbb{E}\left[K\left(\frac{K^{R}}{K}\right)^{2}\right] = \mathbb{E}\left[K^{-1}(K^{R})^{2} \mid K, \rho\right]\right)$$

$$\stackrel{\text{LIE}}{=} \mathbb{E}\left(\mathbb{E}\left[K^{-1}(K^{R})^{2} \mid K, \rho\right]\right)$$

$$\stackrel{\text{EL}}{=} \mathbb{E}\left(K^{-1}\mathbb{E}\left[(K^{R})^{2} \mid K, \rho\right]^{2}\right)$$

$$\stackrel{\text{Jensen}}{\geq} \mathbb{E}\left(K^{-1}\mathbb{E}\left[K^{R} \mid K, \rho\right]^{2}\right)$$

$$\stackrel{\text{BA}}{=} \mathbb{E}\left[K^{-1}\left(K\rho\right)^{2}\right] = \mathbb{E}\left[K\rho^{2}\right]$$

where we use that $\mathbb{R}_+ \ni z \mapsto z^2$ is convex, the conditional Jensen inequality

$$\mathbb{E}\left[K^{R} \mid K, \rho\right]^{2} \leq \mathbb{E}\left[(K^{R})^{2} \mid K, \rho\right],$$

and the binomial assumption $K^R | K, \rho \sim \text{Binomial}(K, \rho)$, whose expectation is $K\rho$.

Remark that the bias appears in the overestimation of $\mathbb{E}[K\rho^2]$, which is also the term through which arises partial identification. The naive technique wrongfully identifies that quantity by $\mathbb{E}[(K^R)^2/K]$. In contrast, our method gives partial identification of $\mathbb{E}[K\rho^2]$ with sharp correct bounds.

Naive estimator Algebraic simplifications from Equation (3.48) gives

$$\pi_{\text{naive}} = 1 - \frac{\mathbb{E}[K]}{2\mathbb{E}[K^D]} + \frac{\mathbb{E}[K] \mathbb{E}\left[(K^R)^2/K\right]}{2\mathbb{E}[K^R] \mathbb{E}[K^D]}.$$

That unknown population quantity only involves expectations of observed variables. Remember that our Assumption 3.1 posits that we observe an i.i.d. sample $(K_j^R, K_j)_{j=1,...,J}$; $K_j^D = K_j - K_j^R$ is observed too. We follow the method of moment, or analogy principle (Manski, 1988), to define a natural estimator of π_{naive} . For any observed random variable V, we replace the expectation $\mathbb{E}[V]$ by the empirical mean $J^{-1} \sum_{j=1}^J V_j$. The J at numerators and denominators cancel out and we obtain the naive estimator

$$\widehat{\pi}_{\text{naive}} := 1 - \frac{\sum_{j=1}^{J} K_j}{2\sum_{j=1}^{J} K_j^D} + \frac{\left(\sum_{j=1}^{J} K_j\right) \left(\sum_{j=1}^{J} \left\{ (K_j^R)^2 / K_j \right\} \right)}{2 \left(\sum_{j=1}^{J} K_j^R \right) \left(\sum_{j=1}^{J} K_j^D \right)}.$$
(3.49)

As regards its proper definition, the estimator $\hat{\pi}_{naive}$ is well-defined with probability approaching one as J tends to infinity. Indeed, $\mathbb{E}[K] > 0$ and $q \in (0, 1)$ implies the probability that the sums in the denominator are null tends to 0. In a pragmatic perspective, if a term at the denominator is null, it means that the sample at hand reports choices made by only one group out of the two. It is therefore fruitless to investigate polarization.

Consistency of $\hat{\pi}_{naive}$ We obtain the consistency of the estimator towards π_{naive} using the Law of Large Numbers (LLN) and the Continuous Mapping Theorem (CMT), see, for instance, Example 2.1 (Classical limit theorems) and Theorem 2.3 (Continuous mapping) of Van der Vaart (2000). To apply the LLN, we simply require a finite first moment for each of the variables we consider. Without loss of generality (see Appendix 3.B.2), we assume $\mathbb{P}(K > 0) = 1$. We can divide by K without worry. By definition of $K, K^R \leq K$ and $K^D \leq K$ almost surely. Thus, we also have $(K^R)^2/K \leq K$. The variables being non-negative, we have the inequalities with absolute values on each side. Therefore, a finite first moment for K is sufficient to obtain finite first moment for all variables whose empirical means appear in $\hat{\pi}_{naive}$, namely K, K^R , K^D , and $(K^R)^2/K$. The condition allows to apply the LLN. For the CMT, remember that in our DGP, $\mathbb{E}[K^R]$ and $\mathbb{E}[K^D]$ are positive, and the fractions are well-defined. We obtain the following consistency result.

Proposition 3.4 (Consistency of $\hat{\pi}_{naive}$ towards π_{naive}). If Assumption 3.1 holds, $\mathbb{P}(K > 0) = 1$, and $\mathbb{E}[|K|] < +\infty$, then $\hat{\pi}_{naive}$ is well-defined with probability approaching one as the sample size J grows to infinity, and $\hat{\pi}_{naive} \xrightarrow{a.s.}{J \to +\infty} \pi_{naive}$.

Remind that $\pi_{\text{naive}} \geq \pi$. So $\hat{\pi}_{\text{naive}}$ tends to overestimate the real level of polarization.

Appendix 3.C Supplements to the application

3.C.1 Processing

The processing used to obtain the bigrams from the Congressional transcripts mimics Gentzkow et al. (2019)'s operations except for two differences. First, we do not use any restrictions based on the number of occurrences, that is, on bigram frequency. Second, we apply a prior spelling correction in our preferred specification.

We itemize below the successive steps from speeches to bigrams. The operations follow the description found in GST, their Online Appendix, and the codebook of the Congressional Record for the 43rd-114th Congresses: Parsed Speeches and Phrase Counts.²⁶

Data All the data we use comes from that Congressional Record and relies on the considerable work achieved by M. Gentzkow, J.M. Shapiro, and M. Taddy to assemble the dataset and match speeches with identified speakers.

The Record provides bigram counts by parties and by speakers. Nonetheless, in order to avoid bigram frequency thresholds, we start from the speeches, as found in the files speeches_###.txt of the Record, and repeat the processing operations performed by GST.²⁷ In the speeches_###.txt files, the raw transcripts are already transformed to separate the text related to speakers' identities from the text of speeches. Furthermore, some formatting operations are done in the latter (see codebook, Section 2.2, p5):

- (i) removing non-speech text (page headers and footers, section titles, parenthetical insertions, votes, and administrative time allotments);
- (ii) removing apostrophes and replacing commas and semicolons with periods;
- (iii) replacing repeated whitespace characters with a single space;
- (iv) removing punctuation, hyphens, periods, and asterisks that separate the speaker's demarcation from the speech;
- (v) removing whitespace leading and trailing the speech.

From texts to bigrams The next stage consists in transforming the speeches' transcripts into counts of bigrams. It comprises several steps.

1. In our specification with correction, we start by applying a spellchecker to the speeches' transcripts. Without correction, we do not perform that operation.

The spellchecker is designed using SymSpell.²⁸ It takes advantage of the three functions, Lookup, LookupCompound, and WordSegmentation, of that spellchecker, and is inspired by manual inspection of the speeches that reveals a default caused, presumably, by the Optical Character Recognition (OCR) digitalization: some words are concatenated together. In our partial identification framework, we are all the more interested in correcting such errors as, otherwise, they would result in single-occurrence phrases, hence broaden our identified set. Indeed, the concatenation is random, and the two, three, four, or more words that happen to be concatenated are likely to appear as such only once. To illustrate that phenomenon, Figure 3.8 shows the longest bigrams found in GST's vocabulary (file vocab.txt in the Congress Record). Those are extreme cases, but, more generally, such concatenations,

²⁶Gentzkow, Matthew, Jesse M. Shapiro, and Matt Taddy. Congressional Record for the 43rd-114th Congresses: Parsed Speeches and Phrase Counts. Palo Alto, CA: Stanford Libraries [distributor], 2018-01-16. https://data.stanford.edu/congress_text.

²⁷Prior to further selection, GST drops rare bigrams: "Bigrams not spoken at least three times in at least one session are removed to ease the computational burden.", codebook, Appendix A, p12.

²⁸https://github.com/wolfgarbe/SymSpell. We use symspellpy, the Python port of SymSpell.

although not massive, seem to occur regularly, notably in the oldest Congress sessions. In addition to other spelling corrections, our spellchecker can segment such concatenated words before constructing bigrams.

Figure 3.8: Illustration of the OCR concatenation errors.

> !	gst_bigram_vocab_sorted	
# /	A tibble: 4,525,243 x 2	
	phrase	nb_characters_with_space
	<chr></chr>	<int></int>
1	gilman pelosic ampbell lowey green wood delauros laught amend	54
2	militaryindustrialscientificcongression complex	47
3	miamijacksonvill chicagoindianapoliscincinnati	46
4	support mcgovernhoekstrapelosimorellajacksonle	46
5	gephardtobeyroehammerschmidtminetashust amend	45
6	harmanhornmccurdysaxtonsprattjohnsontal amend	45
7	kucinichsandersroslehtinendefaziostearn amend	45
8	stupakellsworthpittssmithkapturdahlkemp amend	45
9	chicagoindianapoliscincinnati detroitchicago	44
10	alabamacoosa apalachicolachattahoocheeflint	43

Note: Screenshot from R that shows the longest bigrams (in terms of number of characters) found in GST's vocabulary (file vocab.txt in the Congress Record).

- 2. The speech is coerced to lowercase and is broken into separate words, treating all non-alphanumeric characters as delimiters.
- 3. General English-language stopwords are removed.²⁹
- 4. Remaining words are reduced to their stems using the Porter2 (English) stemming algorithm.³⁰
- 5. The stemmed words are converted to bigrams following their order in the speech.³¹
- 6. The bigrams are converted into counts of bigrams, which undoes the ordering.

Those operations, except the spelling correction 1., copy GST's processing, and we borrow their description from Section 2.2. of the codebook. They provide the counts of occurrences by parties for a set of bigrams. Nonetheless, not all those bigrams are eventually used in the analysis.

From bigrams to vocabulary Indeed, following GST, the third step selects the bigrams considered as "valid". The resulting subset of valid bigrams constitutes the "vocabulary", namely the set of bigrams included in the analysis. The selection stage is performed by flagging bigrams considered as invalid, then "bigrams without a flag are treated as valid vocabulary" (codebook, Appendix A, p15).

- 1. Bigrams with bad syntax are flagged. A bigram has bad syntax if it contains (see codebook, Appendix A, p13):
 - (i) any numbers, symbols, or punctuation;
 - (ii) fewer than five characters, including space;
 - (iii) a one-letter word;
 - (iv) a word beginning with the first three letters of a month.³²
- 2. Bigrams containing the stem of a U.S.-Congress-specific stopword are flagged. Such stopwords come from three sources:

²⁹We use the same source as GST: http://snowball.tartarus.org/algorithms/english/stop.txt.

 $^{^{30}}$ We use the following Python implementation: https://pypi.python.org/pypi/PyStemmer/1.3.0.

³¹For clarification, three successive words "A B C" result in two bigrams: (A, B) and (B, C).

³²To clarify the terminology, a "bigram" (A, B) is made of two "words": A and B.

- (i) the stopwords manually identified by GST and reported in Table 3.1;
- (ii) the names of states;
- (iii) the last names of all congresspeople recorded in the historical source.
- 3. Bigrams recording procedural speech are flagged. These are bigrams that either directly appear in handbooks describing congressional procedures or frequently co-occur with those direct bigrams.³³

Speaker identification We use directly the matching between speakers and speeches, as well as the speaker-level variables, provided by the Congressional Record data in the files ###_SpeakerMap.txt. As GST, we restrict to identified, Republican or Democrat, voting delegates.

Table 3.1: Manually selected U.S.-Congress-specific stopwords identified by GST (reproduces Table 13 of codebook, Appendix A, p13 and Table 4 of the Online Appendix of GST).

absent	adjourn	ask	can	chairman
$\operatorname{committee}$	con	democrat	etc	gentleladies
gentlelady	gentleman	gentlemen	gentlewoman	gentlewomen
hereabout	hereafter	hereat	hereby	herein
hereinafter	hereinbefore	hereinto	hereof	hereon
hereto	heretofore	hereunder	hereunto	hereupon
herewith	month	mr	mrs	nai
nay	none	now	part	per
pro	republican	say	senator	shall
sir	speak	speaker	tell	thank
thereabout	thereafter	the reagainst	thereat	therebefore
there before	thereby	therefor	therefore	therefrom
therein	thereinafter	thereof	thereon	thereto
theretofore	thereunder	thereunto	thereupon	therewith
there with al	today	whereabouts	whereafter	whereas
whereat	whereby	wherefore	wherefrom	wherein
whereinto	whereof	whereon	whereto	whereunder
whereupon	wherever	wherewith	wherewithal	will
yea	yes	yield		

Bigram frequency restrictions GST's specification includes a final step based on the number of occurrences: "Finally, we restrict attention to phrases spoken at least 10 times in at least one session, spoken in at least 10 unique speaker-sessions, and spoken at least 100 times across all sessions." (GST, p1312). We use the same thresholds to obtain "GST's specification" in Figures 3.2 and 3.3 in order to compare with our approach, which avoids such restrictions.

³³The two handbooks used are *Robert's Rules of Order* (1876) and *Riddick's Senate Procedure* (1992). The previous processing steps are performed on those documents' text to obtain stemmed bigrams present in the manuals. Those bigrams are considered to be procedural. Next, GST uses several rules to classify as procedural other bigrams that frequently co-occur with the manuals' bigrams. We refer to the codebook, Appendix A, p13-14, for further details. In practice, to replicate GST's processing, we use the file procedural.txt in the Congressional Record data that lists the words flagged as procedural.

Examples from raw text to valid bigrams We report below some examples of speeches as they appear in the transcripts and their processing into valid bigrams considered in the analysis. They illustrate that the processing operation entails important modifications that may damage the meaning perceived by human readers.

For information, on average across sessions, for a given session, the corpus of text is such that:

- the total number n of occurrences is of the order of several millions (2 to 13) while the vocabulary size J is between 1 and 3 millions (see Figure 3.12);
- the corpus is made of around 250,000 individual speeches;
- close to 40% of those speeches are reduced to an empty set of valid bigrams by the data-processing operations and are removed from the analysis (see an example of such "procedural" speech in Figure 3.9);
- the distribution of the length of speeches (as measured by the number of simple words or tokens they contain, before any processing) shows a large positive skewness: 125 tokens on average but 25 for the median; there is a majority of short to intermediate speeches and a small fraction of (very) long speeches.

Figure 3.9: Example 1 of a speech and the resulting valid bigrams included in the data.



Note: Raw text as found in the U.S. Congress transcripts and resulting valid bigrams included in the analysis following GST data-processing (with or without spelling correction). This particular speech epitomizes "procedural" speech that are entirely dropped from the analysis as they have low to no semantic meaning.

Figure 3.10: Example 2 of a speech and the resulting valid bigrams included in the data.



Note: Raw text as found in the U.S. Congress transcripts and resulting valid bigrams included in the analysis following GST data-processing (with or without spelling correction). This particular speech is an example of a short, rather procedural, speech, yet not reduced to zero valid bigrams and therefore kept in the data included in the analysis.

Figure 3.11: Example 3 of a speech and the resulting valid bigrams included in the data.



Note: Raw text as found in the U.S. Congress transcripts and resulting valid bigrams included in the analysis following GST data-processing (with or without spelling correction). This particular speech is an example of longer speeches, although there can exist much longer speeches.

3.C.2 Additional figures

This subsection regroups various additional figures related to our application to speech polarization in the U.S. Congress. Absent indications, the data used correspond to our preferred dictionary choice as discussed in Section 3.5.2: suppression of procedural bigrams using GST's rules, spelling correction, but no restriction based on bigram frequency.

Some figures show analyses on subsamples of the data. We recall that we consider two subsampling schemes. (i) At the level of speech: we draw a fraction of speeches among the set of all speeches. For a toy example, if the full sample consists of only Examples 2 and 3 above, a 50% subsample drawn at speech level would be one of the two speeches. (ii) At the level of occurrences stratified by speech: within each speech, we draw a fraction of occurrences. Again, if the full sample consists of only Examples 2 and 3 above, a 50% subsample drawn at occurrence level would consist in the union of the random selection of 2 out of the 4 valid bigrams of speech Example 2 with the random selection of 14 out of the 28 valid bigrams of speech Example 3.



Figure 3.12: Number of distinct bigrams and number of occurrences per session.

Note: Number J of distinct bigrams or dictionary size (green) and total number $n := \sum_{j=1}^{J} K_j$ of occurrences (blue) in each Congressional session. The link between the two curves through the evolution over time illustrates Herdan's law. The graph also evidences an overall increase of the length of the corpus over the period 1873-2016.



Figure 3.13: Number of distinct bigrams as a function of number of occurrences.

Note: Number J of distinct bigrams or dictionary size as a function of the total number $n := \sum_{j=1}^{J} K_j$. One point corresponds to one Congressional session, with lighter color blue indicating a more recent session. The figure illustrates Herdan's law and supports our asymptotic: when the corpus gets larger, both the number of occurrences and the size of the dictionary increase.



Figure 3.14: Number of speeches per session.

Note: Number of distinct speeches per session as found in the raw transcripts (in particular, some speeches are reduced to zero bigrams included in the analysis following the data-processing operations). The graph illustrates the variation of the number of speeches over the period 1873-2016. The decrease in recent years combined with the rise in corpus length visible in Figure 3.12 implies that the speeches have become longer. In the hypothesis of positive correlation across occurrences within speech, this suggests a larger magnitude of that correlation overall in recent years.

In this graph, the Congressional sessions instead of years are used; the 43rd session corresponds to years 1873-1874, the 80th one to years 1947-1948, and the 114h to years 2015-2016. The dashed vertical line in session 112th indicates the switch from the bound editions of the transcripts (43rd to 111th Congresses) to the daily editions (available since 97th and used from the 112th; on that point, we follow GST's main specification).

Figure 3.15: Number of tokens (raw words before processing) per speech.



Note: Average and median number of tokens per speech for each session, where the speeches are as found in the raw transcripts, before any processing. In this context, a token is simply any word separated by white spaces. For instance, in Example 2 (Figure 3.10), "United" counts as one token. Tokens are thus distinct from the valid bigrams included in the analysis. Nonetheless, they give an approximation of the length of speeches in terms of number of valid bigrams.



Figure 3.16: Proportion of one-occurrence bigrams per session for different subsamples.

Note: Proportion of bigrams pronounced only once per session for different subsamples, either at speech or occurrence levels. For a given subsampling fraction, the difference between the two schemes (with a larger proportion for occurrence-level subsampling) suggest the presence of positive correlation across occurrences within speech. The parenthesis (1) for subsamples in the legend indicates that the results are presented for one specific subsample draw instead of an average over several draws. Different draws yield the same conclusions, hence this simplification.



Figure 3.17: Total number of occurrences per session for different subsamples.

Note: Total number n of occurrences per session for different subsamples, either at speech or occurrence levels. As expected, they are almost equal for a given subsampling fraction.



Figure 3.18: Estimated lower bound $\hat{\pi}$ per session for different subsamples.

Note: Estimated lower bound $\hat{\pi}$ per session for different subsamples, either at speech or occurrence levels. We remark that the effect of subsampling differs according to the level of subsampling. Subsamples at speech level with a low fraction of the full sample lead to a larger estimate, which is rather unexpected; hence a possible impact of correlation.



Figure 3.19: Estimated upper bound $\hat{\pi}$ per session for different subsamples.

Note: Estimated upper bound $\hat{\pi}$ per session for different subsamples, either at speech or occurrence levels. We remark that the effect of subsampling is similar for the two levels of subsampling, leading to larger estimated upper bounds.



Figure 3.20: Extrapolated estimator $\hat{\pi}_{\text{extrapolated}}$ (cubic) per session for different subsamples.

Note: Extrapolated estimator $\hat{\pi}_{\text{extrapolated}}$ of π (using $\overline{k} = 8$ and r = 3; cubic extrapolation) per session for different subsamples, either at speech or occurrence levels. Except for quite small fractions (5% and 10%) with a speech-level subsampling scheme, our qualitative conclusions based on the extrapolated estimator appear rather robust to subsampling.



Figure 3.21: Extrapolated estimator $\hat{\pi}_{\text{extrapolated}}$ (quadratic) per session for different subsamples.

Note: Extrapolated estimator $\hat{\pi}_{\text{extrapolated}}$ of π (using $\overline{k} = 6$ and r = 2; quadratic extrapolation) per session for different subsamples, either at speech or occurrence levels. The estimates are close to the cubic extrapolation of Figure 3.20, suggesting the exact choice of the extrapolation does not have much influence in our application.

Appendix 3.D Formal link with the model of Gentzkow et al. (2019)

The literature review in the Introduction compared the method developed in Gentzkow, Shapiro and Taddy (2019) (GST) and ours. The main message was the following: the pros and cons of both methods make them more suited to distinct settings. In a nutshell, our method appears more reliable and accessible since it requires only aggregated data. On the other hand, in contexts where it is possible to identify and keep track of the individuals making choices, GST's method incorporates individual covariates more easily, at the expense of a complex numerical estimation however.

This appendix makes a formal link between our data-generating process and the statistical model considered in GST. More precisely, given that the way to include covariates differ between the two methods, we consider a simpler unconditional version of GST's model.

3.D.1 GST's framework

The fundamental distinction between our framework and GST's relates to the statistical unit of interest. In our case, the options (the words, in the sense of the entries of a dictionary) form the units. We observe a vocabulary, or choice set $\mathcal{J} := \{1, \ldots, J\}$, and $J \in \mathbb{N}^*$ is our sample size, that grows to infinity in our asymptotics. In contrast, the basic statistical units, or observations in GST are the occurrences of words (the choices), and the vocabulary (choice set) \mathcal{J} remains fixed. In that perspective, the sample size is the total number of occurrences $n := \sum_{j=1}^{J} K_j$. An observation $i \in \{1, \ldots, n\}$ is an occurrence of the word $W_i \in \mathcal{J}$ pronounced by an individual belonging to group $G_i \in \{\mathbb{R}, \mathbb{D}\}$, \mathbb{R} standing for Republican (more generally, a reference group) and \mathbb{D} for Democrat (the other group). The data-generating process considered by GST is essentially defined by Assumption 3.3.³⁴

Assumption 3.3 (GST). We observe an i.i.d. sample $(W_i, G_i)_{i=1,...,n}$ such that (W_i, G_i) has the same distribution as (W, G), which satisfies: $\mathbb{P}(G = \mathbb{R}) = q \in (0, 1)$, and, for any $g \in \{\mathbb{R}, \mathbb{D}\}$,

 $W \mid G = g \sim \text{Multinomial}\left(1, (p_1^g, \dots, p_J^g)\right),$

where 1 is the number of trials in the multinomial distribution and the event probabilities verify $p_j^g \ge 0$ for any $j \in \{1, \ldots, J\}$ and $\sum_{i=1}^J p_i^g = 1$.

Upstream of that generative process, GST also sets a discrete choice model that specifies the utility earned by a speaker with given characteristics and party affiliation to pronounce each word. In addition to the party G for each occurrence, GST's method needs to record the identity of the speaker to estimate the discrete choice model defined at the level of speakers. The data requirements are thus higher than the aggregated counts of the model defined by Assumption 3.1.

We reuse the notation q because it has the same meaning as $q := \mathbb{E}[K\rho]/\mathbb{E}[K]$ in our model. It is the probability that an arbitrary occurrence, the occurrences being i.i.d., is said by a Republican and not by a Democrat.

For any party $g \in \{R, D\}$ and any word indexed by $j \in \{1, \ldots, J\}$, p_j^g is the probability, conditional on a member of group g saying a word, that the word in question is the j-th

³⁴As explained, we neglect covariates here and present a simplified version of GST's model.

one in the vocabulary. In that sense, p_j^{R} is the reverse of ρ_j . Bayes' theorem connects the two quantities as explained below in Equation (3.51).

In GST's setting, the vocabulary \mathcal{J} is fixed, and the asymptotics is in the number of speakers, hence in the number of occurrences n^{35} . In that framework, the notion of polarization can be formally defined by computing some distance between $(p_j^{\mathrm{R}})_{j=1,\ldots,J}$ and $(p_j^{\mathrm{D}})_{j=1,\ldots,J}$. In the absence of polarization, members of the two parties have the same probabilities to say each word: $p_j^{\mathrm{R}} = p_j^{\mathrm{D}}$ for any $j \in \{1,\ldots,J\}$. On the contrary, a complete polarization means any word is pronounced by one party only: for any $j \in \{1,\ldots,J\}$, $p_j^{\mathrm{R}} = 0$ or $p_j^{\mathrm{D}} = 0$.

3.D.2 Connection between GST's and our model

Our model and polarization index with latent variables In order to connect with GST's model, we introduce an underlying variable λ into our baseline model of Assumption 3.1.

Assumption 3.4 (DGP with latent variable λ). We observe an *i.i.d.* sample $(K_j^R, K_j)_{j=1,...,J}$ such that $(K_j^R, K_j, \lambda_j, \rho_j)$ has the same distribution as (K^R, K, λ, ρ) , which satisfies:

$$\mathbb{E}[K \mid \rho, \lambda] = \lambda > 0, \qquad (3.50)$$
$$K^{R} \mid K, \rho \sim \text{Binomial}(K, \rho).$$

Except for the unobserved latent variable λ , the conditions are similar to our baseline model with i.i.d. sampling of options, a positive expectation for K, and a binomial distribution for K^R conditional on K and ρ . In that sense, Assumption 3.4 defines a subset of models within those of Assumption 3.1. As before, K is the number of occurrences of a word and ρ is the probability, conditional on the occurrence of the word, that is is said by a Republican. Equation (3.50) states that K is mean independent of ρ conditional on λ . Thus, λ represents the propensity of a word to be said, taking into account the value of ρ that is, the fact that the word might be or not a partian expression. Imposing $\mathbb{E}[K \mid \rho, \lambda]$ equal to λ rather than $f(\lambda)$ for any strictly increasing function $f(\cdot)$ is without loss of generality, since λ could be otherwise changed into $f(\lambda)$.

An interesting special case of Assumption 3.4 arises when $K^D := K - K^R$ and K^R are independent Poisson variables with parameters $\lambda^D > 0$ and $\lambda^R > 0$ respectively. Then, classical probability results imply that:

$$\begin{split} & K \mid \lambda^{\mathrm{D}}, \lambda^{\mathrm{R}} \sim \mathrm{Poisson} \Big(\lambda^{\mathrm{D}} + \lambda^{\mathrm{R}} \Big), \\ & K^{R} \mid K, \lambda^{\mathrm{D}}, \lambda^{\mathrm{R}} \sim \mathrm{Binomial} \bigg(K, \frac{\lambda^{\mathrm{R}}}{\lambda^{\mathrm{D}} + \lambda^{\mathrm{R}}} \bigg) \,, \end{split}$$

and Assumption 3.4 holds with $\lambda = \lambda^{\rm D} + \lambda^{\rm R}$ and $\rho = \lambda^{\rm R} / (\lambda^{\rm D} + \lambda^{\rm R})$.

Assumption 3.4 and the law of iterated expectations enable to rewrite our polarization index π defined in Equation (3.2) with the underlying variable λ :

$$\pi = 1 - \frac{\mathbb{E}[\lambda\rho(1-\rho)]}{2\mathbb{E}[\lambda] q(1-q)} = 1 - \frac{\mathbb{E}[\lambda\rho(1-\rho)]}{2\mathbb{E}[\lambda\rho] \left(1 - \frac{\mathbb{E}[\lambda\rho]}{\mathbb{E}[\lambda]}\right)}$$

 $^{^{35}}$ GST specifies their asymptotics in the Appendix, p1334, "All limits are with respect to N", N being "the number of unique speakers".

Link between GST's probabilities and ours The probabilities (p_1^g, \ldots, p_J^g) , $g \in \{R, D\}$, of GST's set-up relates to our conditional probabilities $(\rho_j)_{j=1,\ldots,J}$ by Bayes' rule:

$$\rho_j = \frac{q p_j^{\rm R}}{q p_j^{\rm R} + (1 - q) p_j^{\rm D}}.$$
(3.51)

Indeed, remember that ρ_j is the probability, conditional on the *j*-th word being pronounced, that it is said by Republican, and $q = \mathbb{P}(G = \mathbb{R})$.

GST's polarization index In their original framework, the index considered by GST assumes q = 1/2 and writes

$$\pi_{\rm GST} := \frac{1}{2} \sum_{j=1}^{J} p_j^{\rm R} \rho_j + p_j^{\rm D} (1 - \rho_j)$$

Here, as discussed in the Introduction, we advocate for a more general index that allows for q to differ from 1/2. Therefore, we extend GST's index to

$$\pi_J := 1 + \frac{1}{4q(1-q)} \left[\sum_{j=1}^J q p_j^{\mathrm{R}} \rho_j + (1-q) p_j^{\mathrm{D}} (1-\rho_j) - 1 \right].$$
(3.52)

 π_J simply generalizes π_{GST} and both coincide when q = 1/2.

Link between π_J and π We now clarify the link between the parameter of interest π_J and our index π . First, in line with the meaning of K_j^g in our model, let define the corresponding variables in GST's framework:

$$K_{j,n}^g := \sum_{i=1}^n \mathbb{1} \{ G_i = g, W_i = j \}, \quad g \in \{ \mathbf{R}, \mathbf{D} \},$$

and $K_{j,n} := K_{j,n}^{\mathbb{R}} + K_{j,n}^{\mathbb{D}}$. Then, under Assumption 3.3, we have:

 $K_{j,n}^{\mathrm{R}} | K_{j,n} \sim \mathrm{Binomial}(K_{j,n}, \rho_j).$

Therefore, the conditional model on K_j^R is the same as ours. Note that in GST's set-up formalized by Assumption 3.3, the vocabulary $\mathcal{J} = \{1, \ldots, J\}$ is fixed and the probabilities $(p_j^R)_{j=1,\ldots,J}, (p_j^D)_{j=1,\ldots,J}$, and q are non-stochastic parameters. The $(\rho_j)_{j=1,\ldots,J}$ are thus fixed too whereas they are random variables in our model of Assumption 3.1.

On the other hand, the i.i.d. condition in Assumption 3.3 is not the same as in Assumption 3.4, since the former implies some negative dependence between the J vectors $(K_{j,n}^{\rm D}, K_{j,n}^{\rm R})_{j=1,\dots,J}$. However, that negative dependence becomes asymptotically negligible.

To see this, let $N \sim \text{Poisson}(n)$ be independent of $(G_i, W_i)_{i=1,\dots,n}$. Then, a classical result asserts that $K_{1,N}^{\text{D}}, K_{1,N}^{\text{R}}, \dots, K_{J,N}^{\text{D}}$ and $K_{J,N}^{\text{R}}$ are independent, with $K_{j,N}^{\text{R}} \sim$ Poisson (nqp_j^{R}) and similarly for $K_{j,N}^{\text{D}}$. Note also that, as n grows large, such a random text size becomes close to n, since by the Central Limit Theorem $N - n = O_P(\sqrt{n})$. In other words, introducing just a little randomness in the text size is sufficient to obtain independence between the variables $K_{1,N}^{\text{D}}, K_{1,N}^{\text{R}}, \dots, K_{J,N}^{\text{D}}, K_{J,N}^{\text{R}}$.

As already explained, GST's model and ours differ as regards asymptotics: GST supposes that J is fixed and n, not J, tends to infinity. This explains why they consider

the probabilities $(p_j^{\rm R}, p_j^{\rm D})_{j=1,...,J}$ as fixed rather than random, and why π_J does not coincide with π as defined in Equation (3.2). Nevertheless, we show below that if $(p_j^{\rm R}, p_j^{\rm D})_{j=1,...,J}$ are constructed from random variables and J tends to infinity with n in GST's framework, then the two parameters eventually coincide. Specifically, our equivalence result is based on the following assumption. Note that we underline the dependence in n hereafter by indexing J as well as the probabilities $(p_i^{\rm R}, p_j^{\rm D})_{j=1,...,J}$ and $(\rho_j)_{j=1,...,J}$ by n.

Assumption 3.5. $\lim_{n\to\infty} J_n = +\infty$ and there exist $(\lambda_j^{\mathrm{R}})_{j\geq 1}$ and $(\lambda_j^{\mathrm{D}})_{j\geq 1}$ i.i.d. such that: (i) $\mathbb{P}(\lambda_1^g > 0) = 1$ and $\mathbb{E}[\lambda_1^g] = 1$, for $g \in \{\mathrm{R}, \mathrm{D}\}$; (ii) $p_{j,n}^g = \lambda_j^g / \left[\sum_{\ell=1}^{J_n} \lambda_\ell^g\right]$ for all $n \geq 1$ and $(j,g) \in \{1,\ldots,J_n\} \times \{\mathrm{R},\mathrm{D}\}$.

The asymptotics where the vocabulary size J_n grows to infinity with the length n of the corpus conforms with the way we talk and write in practice, as revealed by Herdan's law (see Remark 3.1 in Section 3.2.1). In other settings, it is relevant as soon as the options display some capacity constraints, that is, can be chosen at most a certain number of times. In such cases, if the total number n of observed choices goes to infinity, so does the number of distinct options. Assumption 3.5 imposes a probabilistic model on the $(p_{j,n}^g)_{j=1,...,J_n}$. The division by $\sum_{\ell=1}^{J_n} \lambda_{\ell}^g$ ensures that $\sum_{j=1}^{J_n} p_{j,n}^g = 1$. Note that as long as $\mathbb{E}[\lambda_1^g] < +\infty$, $\mathbb{E}[\lambda_1^g] = 1$ is a mere normalization: multiplying all the λ_j^g by any constant would not modify the $(p_{j,n}^g)_{j=1,...,J_n}$.

To make the link between π_J and π , let $\lambda_j := q\lambda_j^{\mathrm{R}} + (1-q)\lambda_j^{\mathrm{D}}$, $\rho_{j,\infty} := q\lambda_j^{\mathrm{R}}/\lambda_j$ for any $j \ge 1$ and

$$\pi_{\infty} := 1 + \frac{1}{4q(1-q)} \left\{ \frac{\mathbb{E}[\lambda \left(\rho_{\infty}^2 + (1-\rho_{\infty})^2\right)]}{\mathbb{E}[\lambda]} - 1 \right\} = 1 - \frac{\mathbb{E}[\lambda \rho_{\infty}(1-\rho_{\infty})]}{2 \mathbb{E}[\lambda] q(1-q)}$$

where (λ, ρ_{∞}) has the same distribution as $(\lambda_1, \rho_{1,\infty})$.

Theorem 3.4. Suppose that Assumptions 3.3, 3.4 and 3.5 hold. Then, as n tends to infinity, $\pi_{J_n} \xrightarrow{P} \pi_{\infty}$.

The intuition behind the result is simply that by (3.51) and under Assumption 3.5, $J_n q p_{j,n}^{\rm R}$ and $J_n (1-q) p_{j,n}^{\rm D}$ converge respectively to $\lambda_j \rho_j$ and $\lambda_j (1-\rho_j)$. Then π_J becomes close to

$$1 + \frac{1}{4q(1-q)} \left\{ \frac{1}{J_n} \sum_{j=1}^{J_n} \lambda_j \left[\rho_{j,n}^2 + (1-\rho_{j,n})^2 \right] - 1 \right\}.$$
 (3.53)

In turn, this average converges to π_{∞} . The latter result is not obvious however: as the $(\rho_{j,n}^2)_{j\geq 1}$ are not independent, one cannot apply directly the Law of Large Numbers to (3.53).

Finally, to explicit the link between our index π and the parameter π_J made by Theorem 3.4, remark that simple algebra yields

$$\pi_{\infty} = 1 - \frac{\mathbb{E}[\lambda \rho_{\infty}(1-\rho_{\infty})]}{2 \mathbb{E}[\lambda] q(1-q)}.$$

Chapter 4

Explicit nonasymptotic bounds for the distance to the first-order Edgeworth expansion

This fourth chapter opens the second part of the manuscript concerned with nonasymptotic inference. It is joint work with two co-authors, Alexis Derumigny and Yannick Guyonvarch. The present version is a revised version of an arXiv preprint (arXiv:2101.05780v1).

Chronologically, this work on Edgeworth expansions and Berry-Esseen bounds started after an initial version of Chapter 5 in which we used existing Berry-Esseen inequalities valid under finite third-order moments. However, in the econometrics of linear regressions, it is standard to assume finite *fourth*-order moments for regressors to have a consistent estimator of the asymptotic variance of the OLS estimator and thus classical asymptotic confidence intervals and tests. Refined bounds with as small as possible numeric constants are important for the practical use of nonasymptotic inference tools. Hence the initial motivation of this chapter: improve existing Berry-Esseen bounds under finite fourth-order moments. We do so through bounds for Edgeworth expansions, which basically refines Berry-Essen inequalities by adjusting for possible skewness. The project then extends to study both *i.n.i.d.* and *i.i.d.* cases as well as tighter bounds under additional regularity assumptions, which, in essence, relate to having an absolutely continuous distribution for the observations with respect to Lebesgue's measure (as opposed to a discrete distribution).

Abstract In this article, we study bounds on the uniform distance between the cumulative distribution function of a standardized sum of independent centered random variables with moments of order four and its first-order Edgeworth expansion. Existing bounds are sharpened in two frameworks: when the variables are independent but not identically distributed and in the case of independent and identically distributed random variables. Improvements of these bounds are derived if the third moment of the distribution is zero. We also provide adapted versions of these bounds under additional regularity constraints on the tail behavior of the characteristic function. We finally present an application of our results to the lack of validity of one-sided tests based on the normal approximation of the mean for a fixed sample size.

4.1 Introduction

As the number of observations n in a statistical experiment goes to infinity, many statistics of interest have the property to converge weakly to a $\mathcal{N}(0, 1)$ distribution, once adequately centered and scaled, see, e.g., van der Vaart (2000, Chapter 5) for a thorough introduction. Hence, when little is known on the distribution of a standardized statistic S_n for a fixed sample size n > 0, a classical approach to conduct inference on the parameters of the statistical model amounts to approximate the distribution of that statistic by its tractable Gaussian limit.

A natural and recurring theme in statistics and probability is thus to quantify how far the $\mathcal{N}(0,1)$ distribution lies from the unknown distribution of S_n for a given n. This article aims to present some refined results in one of the simplest and most studied cases: when S_n is a standardized sum of independent random variables. We consider independent but not necessarily identically distributed random variables to encompass a broader range of applications. For instance, certain bootstrap schemes such as the multiplier ones (see Chapter 9 in van der Vaart and Wellner (1996) or Chapter 10 in Kosorok (2006)) boil down to studying a sequence of mutually independent not necessarily identically distributed (i.n.i.d.) random variables conditionally on the initial sample.

More formally, let $(X_i)_{i=1,\dots,n}$ be a sequence of i.n.i.d. random variables satisfying for every $i \in \{1,\dots,n\}$, $\mathbb{E}[X_i] = 0$ and $\gamma_i := \mathbb{E}[X_i^4] < +\infty$. We also define the standard deviation B_n of the sum of the X_i 's, i.e., $B_n := \sqrt{\sum_{i=1}^n \mathbb{E}[X_i^2]}$, so that the standardized sum can be written as $S_n := \sum_{i=1}^n X_i/B_n$. Finally, we use the average individual standard deviation $\overline{B}_n := B_n/\sqrt{n}$ and the average standardized third raw moment $\lambda_{3,n} := \frac{1}{n} \sum_{i=1}^n \mathbb{E}[X_i^3]/\overline{B}_n^3$. The main results of this article are of the form

$$\underbrace{\sup_{x \in \mathbb{R}} \left| \mathbb{P}\left(S_n \le x\right) - \Phi(x) - \frac{\lambda_{3,n}}{6\sqrt{n}} (1 - x^2)\varphi(x) \right|}_{=: \Delta_{n,\mathrm{E}}} \le \delta_n, \tag{4.1}$$

where Φ is the cumulative distribution function of a standard Gaussian random variable, φ is the density function associated with Φ , and δ_n is a positive sequence that depends on the first four moments of $(X_i)_{i=1,...,n}$ and tends to zero under some regularity conditions. In the following, we use the notation $G_n(x) := \Phi(x) + \lambda_{3,n} (6\sqrt{n})^{-1} (1-x^2) \varphi(x)$.

The quantity $G_n(x)$ is usually called the one-term Edgeworth expansion of $\mathbb{P}(S_n \leq x)$, hence the letter E in the notation $\Delta_{n,E}$. Controlling the uniform distance between $\mathbb{P}(S_n \leq x)$ and $G_n(x)$ has a long tradition in statistics and probability, see for instance Esseen (1945) and the books by Cramer (1962) and Bhattacharya and Ranga Rao (1976). As early as in the work of Esseen (1945), it was acknowledged that in independent and identically distributed (*i.i.d.*) cases, $\Delta_{n,E}$ was of the order $n^{-1/2}$ in general and of the order n^{-1} if $(X_i)_{i=1,\dots,n}$ has a nonzero continuous component. These results were then extended in a wide variety of directions, often in connection with bootstrap procedures, see for instance Hall (1992) and Lahiri (2003) for the dependent case.

A one-term Edgeworth expansion can be seen as a refinement of the so-called Berry-Esseen inequality (Berry (1941), Esseen (1942)) which goal is to bound

$$\Delta_{n,\mathrm{B}} := \sup_{x \in \mathbb{R}} |\mathbb{P} \left(S_n \le x \right) - \Phi(x)|$$

The refinement stems from the fact that in $\Delta_{n,E}$, the distance between $\mathbb{P}(S_n \leq x)$ and $\Phi(x)$ is adjusted for the presence of nonasymptotic skewness in the distribution of S_n . Contrary

to the the literature on Edgeworth expansions, there is a substantial amount of work devoted to explicit constants in the Berry-Esseen inequality and its extensions, see, e.g., Bentkus and Götze (1996), Bentkus (2003), Pinelis and Molzon (2016), Chernozhukov et al. (2017), Raič (2018), Raič (2019). The sharpest known result in the *i.n.i.d.* univariate framework is due to Shevtsova (2013), which shows that for every $n \in \mathbb{N}^*$, if $\mathbb{E}[|X_i|^3] < +\infty$ for every $i \in \{1, ..., n\}$, then $\Delta_{n,B} \leq 0.5583 K_{3,n}/\sqrt{n}$ where $K_{p,n} := n^{-1} \sum_{i=1}^{n} \mathbb{E}[|X_i|^p]/(\overline{B}_n)^p$, for $p \in \mathbb{N}^*$, denotes the average standardized *p*-th absolute moment. $K_{p,n}$ measures the tail thickness of the distribution, with $K_{2,n}$ normalized to 1 and $K_{4,n}$ the kurtosis. An analogous result is given in Shevtsova (2013) under the *i.i.d.* assumption where 0.5583 is replaced with 0.4690. A close lower bound is due to Esseen (1956): there exists a distribution such that $\Delta_{n,B} = (C_B/\sqrt{n}) \left(n^{-1} \sum_{i=1}^{n} \mathbb{E}[|X_i|^3]/\overline{B}_n^3\right)$ with $C_B \approx 0.4098$. Another line of research applies Edgeworth expansions in order to get a bound on $\Delta_{n,B}$ that contains higher-order terms, see Adell and Lekuona (2008), Boutsikas (2011) and Zhilova (2020).

Despite the breadth of those theoretical advances, there remain some limits to take full advantage of those results even in simple statistical applications, for instance, when conducting inference on the expectation of a real random variable.¹ If we focus on Berry-Esseen inequalities, Example 4.1 shows that even the sharpest upper bound to date on $\Delta_{n,B}$ can be uninformative when conducting inference on an expectation even for nlarger than 59,000. Therefore, it is natural to wonder whether bounds derived from a one-term Edgeworth expansion could be tighter in moderately large samples (such as a few thousands). In the *i.i.d.* case and under some smoothness conditions, Senatov (2011) obtains such improved bounds. To our knowledge, the question is nevertheless still open in the *i.n.i.d.* setup, as well as in the general setup when no condition on the characteristic function is assumed. In particular, most articles that present results of the form of (4.1) do not provide a fully explicit value for δ_n , that is, δ_n is defined up to some "universal" but unknown constant.

In this article, we derive novel inequalities of the form of (4.1) that aim to be relevant in practical applications. Such "user-friendly" bounds seek to achieve two goals. First, we provide explicit values for δ_n . Second, the bounds δ_n should be small enough to be informative even in small ($n \approx$ hundreds) to moderate ($n \approx$ thousands) sample sizes. We obtain these bounds in an *i.i.d.* setting and in a more general *i.n.i.d.* case only assuming finite fourth moments.

We give improved bounds on $\Delta_{n,E}$ when we assume some regularity assumptions on the tail behavior of the characteristic function f_{S_n} of S_n . Such conditions are related to the continuity of the distribution of S_n with respect to Lebesgue's measure and the differentiability of the corresponding density. These are well-known conditions required for the Edgeworth expansion to be a good approximation of $\mathbb{P}(S_n \leq \cdot)$ with fast rates. Our main results are summed up in Table 4.1.

¹In this article, we only give results for *standardized* sums of random variables, i.e., sums that are rescaled by their standard deviation. In practice, the variance is unknown and has to be replaced with some empirical counterpart, leading to what is usually called a *self-normalized* sum. This is an important question in practice that we leave aside for future research. There exist numerous results on self-normalized sums in the fields of Edgeworth expansions and Berry-Esseen inequalities (Hall (1987), de la Peña et al. (2009)). However, the practical limitations that we point out in this work still prevail.

Setup	General case	Under regularity assumptions on f_{S_n}
i.n.i.d.	$\frac{0.3990K_{3,n}}{\sqrt{n}} + O(n^{-1})$	$\frac{0.195 K_{4,n} + 0.038 \lambda_{3,n}^2}{n} + O(n^{-5/4} + n^{-p/2})$
	(Theorem 4.1)	(Corollary 4.3)
<i>i.i.d.</i>	$\frac{0.1995(K_{3,n}+1)}{\sqrt{n}} + O(n^{-1})$	$\frac{0.195 K_{4,n} + 0.038 \lambda_{3,n}^2}{n} + O(n^{-5/4})$
	(Theorem 4.1)	(Corollary 4.5)

Table 4.1: Summary of the new bounds on $\Delta_{n,E}$ under different scenarios. All the remainder terms are given with explicit expressions for any sample size and are significantly reduced when there is no skewness. For this summary, they are expressed with $O(\cdot)$ for simplicity and to indicate their asymptotic behavior when n goes to infinity. $p \ge 1$ is a constant depending on the smoothness of the characteristic function f_{S_n} .

In the rest of this section, we provide more details about the lack of information given by the Berry-Esseen inequality in Example 4.1 and introduce notation used in the rest of the paper. Section 4.2 presents our bounds on $\Delta_{n,E}$ in the general *i.n.i.d.* case and in the *i.i.d.* setting. In Section 4.3, we develop tighter bounds under regularity assumptions on the characteristic function of S_n . In Section 4.4, we apply our results to show that one-sided tests based on the normal approximation of a sample mean do not hold their nominal level in the presence of nonasymptotic skewness. The main proofs are gathered in Section 4.A and some useful lemmas are proved in Section 4.B.

Example 4.1 (The lack of information conveyed by the Berry-Esseen inequality for inference on an expectation). Let $(Y_i)_{i=1,...,n}$ an i.i.d. sequence of random variable with expectation μ , known variance σ^2 and finite fourth moment with $K_4 := \mathbb{E}\left[(Y_n - \mu)^4\right]/\sigma^4$ the kurtosis of the distribution of Y_n . We want to conduct a test with null hypothesis $H_0: \mu = \mu_0$, for some fixed real number μ_0 , and alternative $H_1: \mu > \mu_0$ with a type-one error at most $\alpha \in (0, 1)$, and ideally equal to α . The classical approach to this problem amounts to comparing $S_n = \sqrt{n}(n^{-1}\sum_{i=1}^n Y_i - \mu_0)/\sigma$ ($X_i := Y_i - \mu_0$) with the $1 - \alpha$ quantile of the $\mathcal{N}(0, 1)$ distribution, $q_{\mathcal{N}(0,1)}(1 - \alpha)$, and reject H_0 if S_n is larger. The Berry-Esseen inequality enables to quantify the mistake caused by using the $\mathcal{N}(0, 1)$ approximation

$$\left|\mathbb{P}\left(S_n \le q_{\mathcal{N}(0,1)}(1-\alpha)\right) - (1-\alpha)\right| \le \frac{0.4690 \,\mathbb{E}\left[|Y_n - \mu_0|^3\right]}{\sqrt{n\sigma^3}} \le \frac{0.4690 \,K_4^{3/4}}{\sqrt{n}}, \qquad (4.2)$$

where the probability and expectation operators are to be understood under the null hypothesis H_0 . Inequality (4.2) is called weakly informative as long as $0.4690K_4^{3/4}/\sqrt{n} \ge \alpha$. Indeed, in that case, we cannot exclude that $\mathbb{P}\left(S_n \le q_{\mathcal{N}(0,1)}(1-\alpha)\right)$ is arbitrarily close to 1, equivalently, that the probability to reject H_0 is arbitrarily close to 0, and therefore that the test is arbitrarily conservative (type-one error arbitrarily small $< \alpha$).

We denote by $n_{\max}(\alpha)$ the largest weakly informative n for a given level α . We note that imposing $K_4 \leq 9$ allows for a wide family of distributions used in practice: any Gaussian, Gumbel, Laplace, Uniform, or Logistic distribution satisfies it, as well as any Student with at least 5 degrees of freedom, any Gamma or Weibull with shape parameter at least 1. Plugging $K_4 = 9$ into (4.2), we remark that for $\alpha = 0.10$, the bound is weakly informative for $n \leq 593$, namely $n_{\max}(0.10) = 593$. For $\alpha = 0.05$, we obtain $n_{\max}(\alpha) = 2,375$. Finally, the situation deteriorates strikingly for $\alpha = 0.01$ where the bound is weakly informative for $n \leq n_{\max}(\alpha) = 59,389$.

Additional notation \lor (resp. \land) denotes the maximum (resp. minimum) operator. For a random variable X, we denote its probability distribution by P_X . For a distribution P, let f_P denote its characteristic function; similarly, for a random variable X, we denote by f_X its characteristic function. We recall that $f_{\mathcal{N}(0,1)}(t) = e^{-t^2/2}$. For two sequences (a_n) , (b_n) , we write $a_n = O(b_n)$ whenever there exists C > 0 such that $a_n \leq Cb_n$; $a_n = o(b_n)$ whenever $a_n/b_n \to 0$; and $a_n \asymp b_n$ whenever $a_n = O(b_n)$ and $b_n = O(a_n)$. We denote by χ_1 the constant $\chi_1 := \sup_{x>0} x^{-3} |\cos(x) - 1 + x^2/2| \approx 0.099$ (Shevtsova, 2010), and by θ_1^* the unique root in $(0, 2\pi)$ of the equation $\theta^2 + 2\theta \sin(\theta) + 6(\cos(\theta) - 1) = 0$. We also define $t_1^* := \theta_1^*/(2\pi) \approx 0.64$ (Shevtsova, 2010). For every $i \in \mathbb{N}^*$, we define the individual standard deviation $\sigma_i := \sqrt{\mathbb{E}[X_i^2]}$. Henceforth, we reason for a fixed arbitrary sample size $n \in \mathbb{N}^*$.

4.2 Control of $\Delta_{n,\mathbf{E}}$ under moment conditions only

We start by introducing two versions of our basic assumptions on the distribution of the variables $(X_i)_{i=1,\dots,n}$.

Assumption 4.1 (Moment conditions in the *i.n.i.d.* framework). $(X_i)_{i=1,...,n}$ are independent and centered random variables such that for every i = 1, ..., n, the fourth raw individual moment $\gamma_i := \mathbb{E}[X_i^4]$ is positive and finite.

Assumption 4.2 (Moment conditions in the *i.i.d.* framework). $(X_i)_{i=1,...,n}$ are *i.i.d.* centered random variables such that the fourth raw moment $\gamma_n := \mathbb{E}[X_n^4]$ is positive and finite.

Assumption 4.2 corresponds to the classical *i.i.d.* sampling with finite fourth moment while Assumption 4.1 is its generalization in the *i.n.i.d.* framework. Those two assumptions primarily ensure that enough moments of $(X_i)_{i=1,...,n}$ exist to build a nonasymptotic upper bound on $\Delta_{n,E}$. In some applications, such as the bootstrap, it is required to consider an array of random variables $(X_{i,n})_{i=1,...,n}$ instead of a sequence. For example, Efron (1979)'s nonparametric bootstrap procedure consists in drawing *n* elements in the random sample $(X_{1,n},...,X_{n,n})$ with replacement. Conditional on $(X_{i,n})_{i=1,...,n}$, the *n* values drawn with replacement can be seen as a sequence of *n i.i.d.* random variables with distribution $\frac{1}{n}\sum_{i=1}^{n} \delta_{\{X_{i,n}\}}$, denoting by $\delta_{\{a\}}$ the Dirac measure at a given point $a \in \mathbb{R}$.

Our results encompass these situations directly. Nonetheless, we do not use the array terminology here as our results hold nonasymptotically, i.e., for any fixed sample size n. We can now state the main result of this section. It is proved in Sections 4.A.3 and 4.A.4.

Theorem 4.1 (Control of the one-term Edgeworth expansion with bounded moments of order four). If Assumption 4.1 or Assumption 4.2 holds, we have the bound

$$\Delta_{n,E} \le \frac{0.1995\,\widetilde{K}_{3,n}}{\sqrt{n}} + \frac{0.031\,\widetilde{K}_{3,n}^2 + 0.195\,K_{4,n} + 0.054\,|\lambda_{3,n}|\widetilde{K}_{3,n} + 0.038\,\lambda_{3,n}^2}{n} + r_{1,n}\,,\tag{4.3}$$
where $\widetilde{K}_{3,n} := K_{3,n} + \frac{1}{n} \sum_{i=1}^{n} \mathbb{E} |X_i| \sigma_i^2 / \overline{B}_n^3$ and $r_{1,n}$ is a remainder term that depends only on $\widetilde{K}_{3,n}$, $K_{3,n}$ $K_{4,n}$ and $\lambda_{3,n}$, defined in Equation (4.22) (resp. in Equation (4.30) under Assumption 4.2).

When the fourth moment is bounded, i.e., $K_{4,n} = O(1)$, we get the rate $r_{1,n} = O(n^{-5/4})$. If furthermore $\mathbb{E}[X_i^3] = 0$ for every i = 1, ..., n, then

- *i.* the average third raw moment is null, i.e., $\lambda_{3,n} = 0$,
- ii. the first-order term in the Edgeworth expansion is null too so that $\Delta_{n,B} = \Delta_{n,E}$,
- iii. the residual term $r_{1,n}$ converges to 0 at the faster rate $n^{-3/2}$ under Assumption 4.1 (resp. at the rate n^{-2} under Assumption 4.2).

Note that it is possible to replace $\widetilde{K}_{3,n}$ by the simpler upper bound $2K_{3,n}$ under Assumption 4.1 (resp. by $K_{3,n} + 1$ under Assumption 4.2). This theorem displays a bound of order $n^{-1/2}$ on $\Delta_{n,E}$. The rate $n^{-1/2}$ cannot be improved when only assuming moment conditions on $(X_i)_{i=1,\dots,n}$ (Esseen (1945), Cramer (1962)). The numerical constants that show up in the leading terms of the bound are quite small. Numerical evaluation of the remainder terms $r_{1,n}$ using its explicit expression shows it is small in practice too. Another nice aspect of those bounds is their dependence on $\lambda_{3,n}$. For many classes of distributions, $\lambda_{3,n}$ can, in fact, be exactly zero. This is the case if for every $i = 1, \dots, n, X_i$ has a non-skewed distribution, such as any distribution that is symmetric around its expectation. More generally, $|\lambda_{3,n}|$ can be substantially smaller than $K_{3,n}$, decreasing the related terms.

This theorem further gives a bound on $\Delta_{n,B}$, even if the third raw moments are not null. Indeed, using Theorem 4.1, the bound $(1 - x^2)\varphi(x)/6 \leq \varphi(0)/6 \leq 0.0665$ for all $x \in \mathbb{R}$, and applying the triangle inequality, we remark that for every $n \in \mathbb{N}^*$

$$\Delta_{n,\mathrm{B}} \le \frac{0.1995 \tilde{K}_{3,n} + 0.0665 |\lambda_{3,n}|}{\sqrt{n}} + O(n^{-1}).$$
(4.4)

Under Assumption 4.1, using the refined inequality $|\lambda_{3,n}| \leq 0.621 K_{3,n}$ (Pinelis, 2011, Theorem 1), we can derive a simpler bound that involves only $K_{3,n}$

$$\frac{0.1995\bar{K}_{3,n} + 0.0665|\lambda_{3,n}|}{\sqrt{n}} \le \frac{0.4403K_{3,n}}{\sqrt{n}}$$

The bound $\Delta_{n,B} \leq 0.4403K_{3,n}/\sqrt{n} + O(n^{-1})$ is already tighter than the sharpest known Berry-Esseen inequality in the *i.n.i.d.* framework, $\Delta_{n,B} \leq 0.5583K_{3,n}/\sqrt{n}$, as soon as the remainder term $O(n^{-1})$ is smaller than the difference $0.118K_{3,n}/\sqrt{n}$. Even in the *i.i.d.* case, this bound is still tighter than the sharpest known Berry-Esseen inequality, $\Delta_{n,B} \leq 0.4690K_{3,n}/\sqrt{n}$, up to a $O(n^{-1})$ term. We refer to Example 4.2 and Figure 4.1 for a numerical comparison, showing improvements for *n* of the order of a few thousands.

The most striking improvement is obtained in the unskewed case when $\mathbb{E}[X_i^3] = 0$ for every integer *i*. In this case, Theorem 4.1 and the inequality $\widetilde{K}_{3,n} \leq 2K_{3,n}$ yield $\Delta_{n,B} \leq 0.3990K_{3,n}/\sqrt{n} + O(n^{-1})$. Note that this result does not contradict Esseen (1956)'s lower bound $0.4098K_{3,n}/\sqrt{n}$ as the distribution he constructs does not satisfy $\mathbb{E}[X_i^3] = 0$ for every *i*.

Under Assumption 4.2, $\widetilde{K}_{3,n} = K_{3,n} + 1$ and we can combine this with (4.4) and the inequality $|\lambda_{3,n}| \leq 0.621 K_{3,n}$ (Pinelis, 2011, Theorem 1), so that we obtain

$$\Delta_{n,\mathrm{B}} \le \frac{0.1995(K_{3,n}+1) + 0.0665 \times 0.621K_{3,n}}{\sqrt{n}} + O(n^{-1}) \le \frac{0.2408K_{3,n} + 0.1995}{\sqrt{n}} + O(n^{-1})$$

One may find this result surprising, given that the numerical constant in front of $K_{3,n}$ in the leading term is smaller than the lower bound constant $C_B := 0.4098$ derived in Esseen (1956). The point is addressed in detail in Shevtsova (2012), where the author explains that the constant cannot be improved if one seeks control of $\Delta_{n,B}$ with a leading term of the form $CK_{3,n}/\sqrt{n}$ for some C > 0. In contrast, our bound on $\Delta_{n,B}$ exhibits a leading term of the form $(CK_{3,n} + c)/\sqrt{n}$ for positive constants C and c.

Example 4.2 (Implementation of our bounds on $\Delta_{n,B}$). Theorem 4.1 provides new tools to control $\Delta_{n,B}$, and we compare them with existing results. To compute our bounds on $\Delta_{n,B}$ as well as previous ones, we need numerical values for $\widetilde{K}_{3,n}$, $\lambda_{3,n}$, and $K_{4,n}$ or upper bounds thereon. A bound on $K_{4,n}$ is in fact sufficient to control $\lambda_{3,n}$ and $K_{3,n}$: Pinelis (2011) ensures $|\lambda_{3,n}| \leq 0.621 K_{3,n}$, and a convexity argument yields $K_{3,n} \leq K_{4,n}^{3/4}$. Moreover, since the third standardized absolute moment has no particular statistical meaning, it is not intuitive to find a natural bound on $K_{3,n}$. On the contrary, the fourth standardized moment $K_{4,n}$ is well-known as the kurtosis of a distribution (the thickness of the tails compared to the central part of the distribution). As explained in Example 4.1, in the i.i.d. framework, imposing $K_{4,n} \leq 9$ is a reasonable assumption. For the sake of comparison, we also impose $K_{4,n} \leq 9$ in the i.n.i.d. case.

We consider the improved bounds that rely on $\lambda_{3,n} = 0$ as well. The different bounds (including the remainder term $r_{1,n}$ for which explicit expressions are given in the proof of Theorem 4.1) are plotted as a function of n in Figure 4.1:

- Shevtsova (2013) i.n.i.d.: $\frac{0.5583}{\sqrt{n}}K_{3,n} \leq \frac{0.5583}{\sqrt{n}}K_{4,n}^{3/4}$
- Shevtsova (2013) i.i.d.: $\frac{0.4690}{\sqrt{n}}K_{3,n} \leq \frac{0.4690}{\sqrt{n}}K_{4,n}^{3/4}$
- Theorem 4.1 i.n.i.d.: $\frac{0.4403}{\sqrt{n}}K_{3,n} + r_{1,n} \le \frac{0.4403}{\sqrt{n}}K_{4,n}^{3/4} + r_{1,n}$
- Theorem 4.1 i.n.i.d. (unskewed): $\frac{0.3990}{\sqrt{n}}K_{3,n}^3 + r_{1,n} \leq \frac{0.3990}{\sqrt{n}}K_{4,n}^{3/4} + r_{1,n}$

$$- Theorem 4.1 \ i.i.d.: \ \frac{0.2408K_{3,n}+0.1995}{\sqrt{n}} + r_{1,n} \le \frac{0.2408K_{4,n}^{3/4}+0.1995}{\sqrt{n}} + r_{1,n}$$

 $\begin{array}{l} - \ Theorem \ 4.1 \ i.i.d. \ (unskewed): \ \frac{0.1995K_{3,n}+0.1995}{\sqrt{n}} + r_{1,n} \leq \frac{0.1995K_{4,n}^{3/4}+0.1995}{\sqrt{n}} + r_{1,n} \\ As \ previously \ mentioned, \ our \ bound \ in \ the \ baseline \ i.n.i.d. \ case \ gets \ close \ to \ the \ \ the \$

As previously mentioned, our bound in the baseline i.n.i.d. case gets close to the best known Berry-Esseen bound in the i.i.d. setup (Shevtsova, 2013) when n is larger than 10,000. When $\lambda_{3,n} = 0$, our bounds are smaller, notably in the i.i.d. scenario: in the latter case, the bound is smaller than 0.05 for $n \approx 1,000$, highlighting that taking (lack of) skewness into account matters in improving Berry-Esseen bounds. More generally, the results are considerably better in the i.i.d. framework.

4.3 Improved bounds on $\Delta_{n,\mathbf{E}}$ under assumptions on the tail behavior of f_{S_n}

In this section, we derive tighter bounds on $\Delta_{n,E}$ under additional regularity conditions on the tail behavior of the characteristic function of S_n . They follow from Theorem 4.2, which provides an alternative upper bound on $\Delta_{n,E}$ that involves the tail behavior of f_{S_n} . This theorem is proved in Section 4.A.5 in the *i.n.i.d.* setting under Assumption 4.1 (resp. in Section 4.A.6 in the *i.i.d.* case with Assumption 4.2).

Theorem 4.2. If Assumption 4.1 or Assumption 4.2 holds, we have the bound

$$\Delta_{n,E} \le \frac{0.195 K_{4,n} + 0.038 \lambda_{3,n}^2}{n} + \frac{1.0253}{\pi} \int_{a_n}^{b_n} \frac{|f_{S_n}(t)|}{t} dt + r_{2,n}, \qquad (4.5)$$



Figure 4.1: Comparison between existing (Shevtsova, 2013) and new (Theorem 4.1) Berry-Esseen upper bounds on $\Delta_{n,B} := \sup_{x \in \mathbb{R}} |\mathbb{P}(S_n \leq x) - \Phi(x)|$ for different sample sizes under moment conditions only (log-log scale). Note that, compared to existing ones, the new bounds make use of the assumption of finite fourth moment.

where $a_n := 2t_1^* \pi \sqrt{n} / \widetilde{K}_{3,n}$, $b_n := 16\pi^4 n^2 / \widetilde{K}_{3,n}^4$, and $r_{2,n}$ is a remainder term that depends only on $\widetilde{K}_{3,n}$, $K_{3,n}$, $K_{4,n}$ and $\lambda_{3,n}$. The term $r_{2,n}$ is defined in Equation (4.35) under Assumption 4.1 (resp. in Equation (4.40) under Assumption 4.2). When the fourth moment is bounded, i.e., $K_{4,n} = O(1)$, we have $r_{2,n} = O(n^{-5/4})$.

- If furthermore $\mathbb{E}[X_i^3] = 0$ for every i = 1, ..., n, then
 - *i.* the average third raw moment is null, i.e., $\lambda_{3,n} = 0$,
 - ii. the first-order term in the Edgeworth expansion is null too so that $\Delta_{n,B} = \Delta_{n,E}$,
 - iii. the residual term $r_{2,n}$ converges to 0 at the faster rate $n^{-3/2}$ under Assumption 4.1 (resp. at the rate n^{-2} under Assumption 4.2).

This theorem is obtained under the same conditions as Theorem 4.1. The first term contains quantities that were already present in the term of order 1/n in the bound of Theorem 4.1: $0.195K_{4,n}$ and $0.038\lambda_{3,n}^2$. On the contrary, the other terms are encompassed in the integral term and in the remainder. Indeed, a careful reading of the proofs (see notably Section 4.A.2 that outlines the structure of the proofs of all theorems) shows that the leading term $0.1995 \widetilde{K}_{3,n}/\sqrt{n}$ in the bound (4.3) comes from choosing a free tuning parameter T of the order of \sqrt{n} . Here, we make another choice for T such that this term is now negligible. The cost of this change of T is the introduction of the integral term involving f_{S_n} . The leading term of the bound thus depends on the tail behavior of f_{S_n} .

The following corollaries specify the magnitude of the bound according to the assumed

regularity conditions on f_{S_n} . The first one considers a polynomial decrease in the *i.n.i.d.* framework.

Corollary 4.3. If Assumption 4.1 holds and for all $t \in (a_n, +\infty)$, $|f_{S_n}(t)| \leq C_0 t^{-p}$ for some constants $C_0, p > 0$, then

$$\Delta_{n,E} \le \frac{0.195 \, K_{4,n} + 0.038 \, \lambda_{3,n}^2}{n} + \frac{1.0253 \, C_0 a_n^{-p}}{\pi} + r_{3,n}$$

where $r_{3,n} := r_{2,n} - 1.0253 C_0 b_n^{-p} / \pi$ and $r_{2,n}$ is defined in Equation (4.35).

Besides moment conditions, Corollary 4.3 requires a uniform control on the tail of f_{S_n} beyond the point $2t_1^*\pi\sqrt{n}/\widetilde{K}_{3,n}$. Whenever $\widetilde{K}_{3,n} = o(\sqrt{n})$, this condition is a tail control of the characteristic function of S_n above a point that tends to infinity, thus making this condition weaker to impose. When p > 1, f_{S_n} is absolutely integrable and thus P_{S_n} is a continuous distribution (Ushakov, 2011, Theorem 1.2.6). In this case, we give a characterization of the tail constraint on f_{S_n} in terms of smoothness of the underlying distribution function in Proposition 4.4.

Although Corollary 4.3 is valid for every positive p, it is only an improvement on the results of the previous section under the stricter condition p > 1. In particular when p = 2, a_n^{-p} is exactly of the order n^{-1} , we obtain

$$\Delta_{n,\mathrm{E}} \le \frac{0.195 \, K_{4,n} + 0.038 \, \lambda_{3,n}^2 + 1.0253 \, C_0 \pi^{-1}}{n} + o(n^{-1}),$$

and when p > 2

$$\Delta_{n,\mathrm{E}} \le \frac{0.195 \, K_{4,n} + 0.038 \, \lambda_{3,n}^2}{n} + o(n^{-1}).$$

Combining these bounds on $\Delta_{n,E}$ with the expression of the Edgeworth expansion translates into upper bounds on $\Delta_{n,B}$ of the form

$$\Delta_{n,\mathrm{B}} \le \frac{0.0665 \,|\lambda_{3,n}|}{\sqrt{n}} + O(n^{-1}) \le \frac{0.0413 \,K_{3,n}}{\sqrt{n}} + O(n^{-1}).$$

In the regime when the $O(n^{-1})$ term is smaller than $0.0413K_{3,n}/\sqrt{n}$, the bound on $\Delta_{n,B}$ becomes much better than $0.5583K_{3,n}/\sqrt{n}$ or $0.4690K_{3,n}/\sqrt{n}$. This can happen even for sample sizes n of the order of a few thousands, assuming that $K_{3,n}$ and $K_{4,n}$ are reasonable. When $\mathbb{E}[X_i^3] = 0$ for every $i = 1, \ldots, n$, we remark that $\Delta_{n,B} = \Delta_{n,E}$, meaning that we obtain a bound on $\Delta_{n,B}$ of order n^{-1} .

We verify these rates through a numerical application in Example 4.6 for the specific choice $C_0 = 1$ and p = 2. These choices are satisfied for usual distributions such as the Laplace distribution (for which these values of C_0 and p are sharp) and the Gaussian distribution.

More generally, a bound on the tail of the characteristic function is nearly equivalent to a regularity condition on the density. We detail this in the following proposition. The first part of this proposition is taken from (Ushakov, 2011, Theorem 2.5.4) (see also Ushakov and Ushakov (1999)). The second part is proved in Section 4.B.5.

Proposition 4.4. Let $p \ge 1$ be an integer, Q be a probability measure that admits a density q with respect to Lebesgue's measure, and f_Q its corresponding characteristic function.

1. If q is (p-1) times differentiable and $q^{(p-1)}$ is a function with bounded variation, then

$$|f_Q(t)| \le \frac{\operatorname{Vari}[q^{(p-1)}]}{|t|^p},$$

where $\operatorname{Vari}[\psi]$ denotes the total variation of a function ψ . 2. If $t \mapsto |t|^{p-1} |f_Q(t)|$ is integrable on a neighborhood of $+\infty$, then q is (p-1) times differentiable.

It is sufficient that there exists C > 0 and $\beta > 1$ such that $|f_Q(t)| \leq C/(|t|^p \log(|t|)^\beta)$ to satisfy the integrability condition in the second part of Proposition 4.4. Proposition 4.4 shows that the tail condition on f_{S_n} in Corollary 4.3 is satisfied if P_{S_n} has a density g_{S_n} with respect to Lebesgue's measure that is p-1 times differentiable and such that its (p-1)-th derivative is of bounded variation with total variation $V_n := \operatorname{Vari}[g_{S_n}^{(p-1)}]$ uniformly bounded in n. In such cases, we can take $C_0 = 1 \vee \sup_{n \in \mathbb{N}^*} V_n$.

Another possibility would be to impose $|f_{S_n}(t)| \leq \max_{1 \leq r \leq M} |\rho_r(t)|$ for every $|t| \geq a_n$ and for $(\rho_r)_{r=1,\dots,M}$ a family of known characteristic functions. Indeed, in a statistical framework, the characteristic function f_{S_n} is unknown. This second suggestion boils down to a semiparametric assumption on P_{S_n} : f_{S_n} is assumed to be controlled in a neighborhood of $+\infty$ by the behavior of at least one of the M characteristic functions $(\rho_r)_{r=1,\dots,M}$, but f_{S_n} needs not be exactly one of those M characteristic functions. This semiparametric restriction becomes less and less stringent as n increases since we need to control f_{S_n} on a region that vanishes as n goes to infinity. Since S_n is centered and of variance 1 by definition, the choice of possible ρ_r is naturally restricted to the set of characteristic functions that correspond to such standardized distributions.

We state a second corollary that deals with the *i.i.d.* framework. We need to define the following quantity $\kappa_n(\widetilde{K}_{3,n}) := \sup_{t: |t| \ge 2t_1^* \pi/\widetilde{K}_{3,n}} |f_{X_n/\sigma_n}(t)|$. Under Assumption 4.2, we remark that $\sup_{t: |t| \ge a_n} |f_{S_n}(t)| = \kappa_n(\widetilde{K}_{3,n})^n$.

Corollary 4.5. If Assumption 4.2 holds and P_{X_n/σ_n} has an absolutely continuous component with respect to Lebesgue's measure, then $\kappa_n(\widetilde{K}_{3,n}) < 1$ and

$$\Delta_{n,E} \le \frac{0.195 \, K_{4,n} + 0.038 \, \lambda_{3,n}^2}{n} + \frac{1.0253 \, \kappa_n(\widetilde{K}_{3,n})^n \log(c_n)}{\pi} + r_{2,n} \,,$$

where $c_n := b_n/a_n = 8\pi^3 n^{3/2}/(t_1^* \widetilde{K}_{3,n}^3)$ and $r_{2,n}$ is defined in Equation (4.40).

Note that for a given s > 0 and a variable Z, $\sup_{t:|t| \ge s} |f_Z(t)| = 1$ if and only if P_Z is a lattice distribution, i.e., concentrated on a set of the form $\{a + nh, n \in \mathbb{Z}\}$ (Ushakov, 2011, Theorem 1.1.3). Therefore, $\kappa_n(\widetilde{K}_{3,n}) < 1$ as soon as the distribution is not lattice, which is the case for any distribution with an absolute continuous component.

In Corollary 4.5, we derive an upper bound of order n^{-1} on $\Delta_{n,E}$. As the first term on the right-hand side is independent of the behavior of the characteristic function, it is unchanged compared to Theorem 4.2 and Corollary 4.3. The second term in the bound, $(1.0253/\pi)\kappa_n(\widetilde{K}_{3,n})^n \log(c_n)$, corresponds to an upper bound on the integral term of Equation (4.5) in Theorem 4.2. In the *i.i.d.* setting with a distribution of the $(X_i)_{i=1,\dots,n}$ independent of n (the case of a sequence of variables as opposed to an array), $\kappa_n(\widetilde{K}_{3,n}) = \kappa$, a constant independent of n. Consequently, since $0 \leq \kappa < 1$ is at the power n, this bound on the integral term of (4.5) is exponentially better than the one derived in Corollary 4.3 and is asymptotically negligible in front of the first term of order $O(n^{-1})$.

In applications, we may want to compute this bound explicitly, so we would need an explicit value for $\kappa_n(K_{3,n})$. This value depends on the (unknown) distribution P_{X_n/σ_n} and if it is too close to one for a given n, $\kappa_n(\widetilde{K}_{3,n})^n$ will not be small. For instance, if n = 1,000 and $\kappa_n(\widetilde{K}_{3,n}) = 0.999$, we have $\kappa_n(\widetilde{K}_{3,n})^n \approx 0.37$. As a result, we must impose that $\kappa_n(K_{3,n})$ be bounded away from one in order to use such a bound in practice, which amounts to placing a restriction on P_{X_n/σ_n} .

This is, in fact, more than a restriction on the tail of f_{X_n/σ_n} since control of the tail of a characteristic function induces control on its central part as well (Ushakov, 2011, Theorem 1.4.4). By convexity $K_{3,n} \geq 1$ and in the most favorable case, we need a control of f_{X_n/σ_n} above the point $2t_1^*\pi$. As in the *i.n.i.d.* case, a possibility to do so is to impose that the characteristic function f_{X_n/σ_n} is controlled by some known family of characteristic functions ρ_1, \ldots, ρ_M beyond $2t_1^*\pi/K_{3,n}$.

In Example 4.6, we derive a bound on $\Delta_{n,B}$ from Corollary 4.5 which has the same flavor as the one obtained from Corollary 4.3. Besides, similarly to what has been done previously, we derive a better result under the assumption that $\mathbb{E}[X_n^3] = 0$.

Example 4.6 (Implementation of our bounds on $\Delta_{n,B}$). We compare the bounds on $\Delta_{n,B}$) obtained in Corollaries 4.3 and 4.5 to $0.5583K_{3,n}/\sqrt{n}$ and $0.4690K_{3,n}/\sqrt{n}$. As in previous examples, we fix $K_{4,n} \leq 9$.

- $Corollary 4.3 i.n.i.d., p = 2 and C_0 = 1: \Delta_{n,B} \leq \frac{0.0413 K_{4,n}^{3/4}}{\sqrt{n}} + \frac{0.195 K_{4,n} + 0.024 K_{4,n}^{3/2}}{n} + \frac{0.018 K_{4,n} + 0.024 K_{4,n}^{3/2}{n} + \frac{0.018 K_{4,n} + 0.024 K_{4,n}$ $1.0253 \, \pi^{-1} a_n^{-2} + r_{3.n}$
- Corollary 4.3 i.n.i.d. unskewed, p = 2 and $C_0 = 1$: $\Delta_{n,B} \le \frac{0.195 K_{4,n}}{n} + 1.0253 \pi^{-1} a_n^{-2} + 1.02$ $r_{3,n}$
- $\ Corollary \ 4.5 \ i.i.d., \ \kappa = 0.99 \colon \Delta_{n,B} \le \frac{0.0413 \, K_{4,n}^{3/4}}{\sqrt{n}} + \frac{0.195 \, K_{4,n} + 0.024 \, K_{4,n}^{3/2}}{n} + \frac{1.0253 \, \kappa^n \log(c_n)}{\pi} + \frac{1.0$ $r_{2,n}$

 $- Corollary 4.5 i.i.d. unskewed, \kappa = 0.99: \Delta_{n,B} \leq \frac{0.195 K_{4,n}}{n} + \frac{1.0253 \kappa^n \log(c_n)}{\pi} + r_{2,n}$ As a reminder, we underscore that the terms $r_{3,n} := r_{2,n} - 1.0253 C_0 b_n^{-p} / \pi$ and $r_{2,n}$ have explicit expressions (Equation (4.35) in the i.n.i.d. case and Equation (4.40) in the i.i.d. case) for any sample size n that we compute to compare the different bounds.

Since we assume some supplementary regularity conditions in this section, it should not be surprising that they lead to an improvement of the upper bounds on $\Delta_{n,B}$. Figure 4.2 displays the different bounds that we obtained as a function of the sample size n, alongside with the existing bounds (Shevtsova, 2013) that do not assume such regularity conditions. The new bounds take advantage of these regularity conditions and are therefore much tighter, especially in the unskewed case where the rate of convergence gets faster from $1/\sqrt{n}$ to 1/n.

Conclusion and statistical applications 4.4

In addition to Figures 4.1 and 4.2, we can illustrate the improvements provided by Theorem 4.1 and Corollary 4.5 in the *i.i.d.* case by resuming Example 4.1. This example considers the classical problem of conducting inference on a univariate expectation from *i.i.d.* observations with known variance. For usual confidence levels and different assumptions, Table 4.2 indicates the largest weakly informative sample size $n_{\max}(\alpha)$ as defined in Example 4.1: for $n \leq n_{\max}(\alpha)$, the Berry-Esseen bound that controls the approximation of $\mathbb{P}(S_n \leq q_{\mathcal{N}(0,1)}(1-\alpha))$ by $\Phi(q_{\mathcal{N}(0,1)}(1-\alpha)) = 1-\alpha$ cannot exclude arbitrarily conservative



Figure 4.2: Comparison between existing (Shevtsova, 2013) and new (Corollaries 4.3 and 4.5) Berry-Esseen upper bounds on $\Delta_{n,B} := \sup_{x \in \mathbb{R}} |\mathbb{P}(S_n \leq x) - \Phi(x)|$ for different sample sizes with additional regularity assumption on f_{S_n} (log-log scale). Note that, compared to existing ones, the new bounds make use of the regularity assumption.

tests (with a probability to reject the null hypothesis arbitrarily close to 0). For these numerical applications, we keep the same constants as in the previous figures, namely $K_{4,n} = 9$ and $\kappa = 0.99$.

As seen in the previous examples, explicit values or bounds on some functionals of P_{S_n} , such as $\lambda_{3,n}$, $\widetilde{K}_{3,n}$, $K_{4,n}$, C_0 , p, and κ , are required to obtain our nonasymptotic bounds on a standardized sample mean. The bound $K_{4,n} \leq 9$ encompasses a wide range of standard distributions and provides a bound on $\lambda_{3,n}$ and $\widetilde{K}_{3,n}$ too. Nonetheless, we may prefer to avoid such a restriction, either because it is unlikely to be satisfied or, on the contrary, because it might be overly conservative. An alternative would be to estimate the moments $\lambda_{3,n}$, $K_{3,n}$ (then $\widetilde{K}_{3,n}$ is upper bounded by $2K_{3,n}$ in *i.n.i.d.* settings or $1 + K_{3,n}$ in *i.i.d.* settings), $K_{4,n}$ using the data. In the *i.i.d.* case, we suggest estimating them by their empirical counterparts (method of moments estimation). We could then compute our bounds by replacing the unknown needed quantities with their estimates. We acknowledge that this type of "plug-in" approach is only approximately valid.

Theorem 4.2, which underlies Corollaries 4.3 and 4.5, involves the integral $\int_{a_n}^{b_n} f_{S_n}(t)/t dt$, which depends on the a priori unknown characteristic function of S_n . A nonparametric possibility would be to estimate f_{S_n} and then to numerically integrate it. A classical estimator is the empirical characteristic function (see for instance (Ushakov, 2011, Chapter 3)). However, this nonparametric estimator may be unstable since the uniform convergence of the empirical characteristic function only holds on compact sets. Therefore, it might

Bound on $\Delta_{n,B}$	$\alpha = 0.10$	$\alpha = 0.05$	$\alpha = 0.01$
Existing	593	2,375	$59,\!389$
Thm. 4.1	462	1,412	$22,\!265$
Thm. 4.1 unskewed	239	764	$15,\!875$
Cor. 4.5	384	497	1,912
Cor. 4.5 unskewed	349	427	636

Table 4.2: In the test of $H_0: \mu = \mu_0$ against $H_1: \mu > \mu_0$ with μ the expectation of *i.i.d.* real random variables with known variance, largest weakly informative sample sizes $n_{\max}(\alpha)$, with $\alpha \in \{0.10, 0.05, 0.01\}$ for different assumptions and Berry-Esseen bounds: existing bound with finite third moment (Shevtsova, 2013), new bound with finite fourth moment (Theorem 4.1, skewed and unskewed cases), new bound with additional regularity condition on f_{X_n/σ_n} ; essentially, a continuous distribution (Corollary 4.5, skewed and unskewed cases).

be of interest to avoid such an estimation. Corollaries 4.3 and 4.5 offer this alternative under some regularity conditions on f_{S_n} . They require an upper control on the tail of the characteristic function. Such a bound can be given using expert knowledge of the regularity of the density of the S_n . The plot of the module of the empirical characteristic function can also be used to back such an intuition.

We now examine some implications of our theoretical results for the nonasymptotic validity of statistical tests based on the Gaussian approximation of the distribution of a sample mean using *i.i.d.* data. As an illustration, we continue Example 4.1. Remember that we observe *i.i.d.* observations $(Y_i)_{i=1,...,n}$ with finite fourth moment. We assume their variance is known and want to conduct inference on the expectation of Y_n . For fixed number $\mu_0 \in \mathbb{R}$ and nominal control $\alpha \in (0, 1)$ on the type-one error, we study the unilateral test $H_0 : \mathbb{E}[Y_n] \leq \mu_0$ against $H_1 : \mathbb{E}[Y_n] > \mu_0$.

The relevant test statistic for this problem is $S_n = \sqrt{n}(n^{-1}\sum_{i=1}^n Y_i - \mu_0)/\sigma$ (that is, the standardized sum S_n with $X_i := Y_i - \mu_0$) and we consider the rejection region $\{S_n > q_{\mathcal{N}(0,1)}(1-\alpha)\}$. In Example 4.1, the probability operator is understood under the null hypothesis. Here, H_0 is composite, and the following probability operators are to be understood for $\mathbb{E}[Y_n] = \mu_0$, which is the most difficult data-generating process within the null to distinguish between the two hypotheses H_0 and H_1 , and thus the one to consider to study the type-one error and p-value of the test.

Starting from Theorem 4.1 or Corollary 4.5, we set $x = q_{\mathcal{N}(0,1)}(1-\alpha)$ (henceforth denoted $q_{1-\alpha}$ to lighten notation) and deduce the following upper and lower bounds

$$\frac{\lambda_{3,n}}{6\sqrt{n}}(1-q_{1-\alpha}^2)\varphi(q_{1-\alpha}) - \delta_n \le \mathbb{P}\left(S_n \le q_{1-\alpha}\right) - (1-\alpha) \le \frac{\lambda_{3,n}}{6\sqrt{n}}(1-q_{1-\alpha}^2)\varphi(q_{1-\alpha}) + \delta_n,$$
(4.6)

where δ_n is the corresponding bound on $\Delta_{n,E}$. We recall that either $\delta_n = O(n^{-1/2})$ under moments conditions only (Section 4.2) or $\delta_n = O(n^{-1})$ under additional regularity conditions (Section 4.3). For the sake of conciseness, we consider here the latter framework where the distribution of Y_n , hence of X_n , has an absolutely continuous component. In that smooth case, since $\delta_n = O(n^{-1})$, we can write, up to the term δ_n ,

$$\mathbb{P}\left(S_n \le q_{1-\alpha}\right) \approx 1 - \alpha + \frac{\lambda_{3,n}}{6\sqrt{n}} (1 - q_{1-\alpha}^2)\varphi(q_{1-\alpha}).$$
(4.7)

As a consequence, for $\alpha \leq 0.15$ (so that $q_{1-\alpha} > 1$), as soon as the skewness $\lambda_{3,n}$ is higher than $\frac{6\sqrt{n}\,\delta_n}{(q_{1-\alpha}^2-1)\varphi(q_{1-\alpha})} > 0$, the probability $\mathbb{P}\left(S_n \leq q_{1-\alpha}\right)$ has to be smaller than $1-\alpha$. In that case, since the event $\{S_n \leq q_{1-\alpha}\}$ is the complementary of the rejection region, the probability of rejecting H_0 under the null exceeds α ; in other words, the test does not reach its stated nominal control α on the type-one error and is said liberal.

Conversely, when $\lambda_{3,n}$ is lower than $\frac{-6\sqrt{n}\,\delta_n}{(q_{1-\alpha}^2-1)\varphi(q_{1-\alpha})} < 0$, the probability $\mathbb{P}\left(S_n \leq q_{1-\alpha}\right)$ has to be larger than $1-\alpha$; equivalently, the probability to reject under the null is below α . In that case, the test is said conservative.

The distortion can also be seen in terms of p-values. In the unilateral test we consider, the p-value is $pval := \mathbb{P}(S_n > s_n)$ with s_n the observed value of S_n in the sample. In contrast, the approximated p-value is $\widetilde{pval} := 1 - \Phi(s_n)$. Analogous to Equation (4.6), setting $x = s_n$ yields

$$\frac{\lambda_{3,n}}{6\sqrt{n}}(1-s_n^2)\varphi(s_n) - \delta_n \le (1-pval) - (1-\widetilde{pval}) \le \frac{\lambda_{3,n}}{6\sqrt{n}}(1-s_n^2)\varphi(s_n) + \delta_n.$$

Therefore,

$$pval - \frac{\lambda_{3,n}}{6\sqrt{n}}(s_n^2 - 1)\varphi(s_n) - \delta_n \le \widetilde{pval} \le pval - \frac{\lambda_{3,n}}{6\sqrt{n}}(s_n^2 - 1)\varphi(s_n) + \delta_n$$

When P_{Y_n} has an absolutely continuous component, $\delta_n = O(n^{-1})$ by Corollary 4.5 and negligible compared to the skewness term involving $\lambda_{3,n}$. This bound indicates that for $|s_n| > 1$, the higher the skewness, the smaller the approximated p-value compared to the true one. This results in overconfidence in rejecting the null. The distortion towards lower approximated p-values occurs in a similar way whether $s_n < -1$ or $s_n > 1$ due to the parity of the function $x \mapsto (x^2 - 1)\varphi(x)$. Conversely, negative skewness leads to the under-rejection of the null hypothesis. Table 4.3 describes the sign of the distortion according to the skewness of P_{Y_n} (the sign of $\lambda_{3,n}$) and the observed value s_n of the test statistic.

More generally, for any n and point x, Equation (4.6) shows that $\mathbb{P}(S_n \leq x)$ belongs to the interval $[\Phi(x) + \lambda_{3,n}(1-x^2)\varphi(x)/(6\sqrt{n}) \pm \delta_n]$, which is not centered at $\Phi(x)$ whenever $\lambda_{3,n} \neq 0$. Figures 4.3 and 4.4 illustrate this distortion in terms of p-values for different sample sizes and observed test statistics in the case of a positive skewness. The length of the interval does not depend on x and shrinks at speed δ_n . On the contrary, the location of the interval depends on x. For given skewness $\lambda_{3,n}$ and sample size n, it is all the more shifted away from the asymptotic approximation $\Phi(x)$ as $x \mapsto (1-x^2)\varphi(x)$ is large in absolute value. That function has a global maximum at x = 0 and minima at the points $-x^*$ and $x^* \approx 1.73$ (Figure 4.5 plots the function for illustration). Consequently, irrespective of n, the largest gaps between $\mathbb{P}(S_n \leq x)$ and $\Phi(x)$ may be expected around

Skewness	$s_n < -1$	$-1 < s_n < 1$	$s_n > 1$
Negative	+	—	+
Positive	_	+	—

Table 4.3: Sign of the distortion (approximated p-value) pval - pval (true p-value) in a one-sided test on the expectation of *i.i.d.* and continuous observations.



x = 0 or $x = \pm x^*$. Inversely, when $x = \pm 1$, the gap is null and the interval is exactly centered at $\Phi(x)$.

Figure 4.3: For every possibly observed test statistic $s \in \mathbb{R}$, we represent the possible pairs (pval(s), pval(s)) (shaded region) of the true p-value $pval(s) := 1 - \mathbb{P}(S_n \leq s)$ and the approximated one $pval(s) := 1 - \Phi(s)$. The boundaries of this region correspond to the curves $(1 - \Phi(s), 1 - \Phi(s) + (\lambda_{3,n}/6\sqrt{n})(s^2 - 1)\varphi(s) \pm \delta_n)_{s \in \mathbb{R}}$ where δ_n is the bound on $\Delta_{n,E}$ given by Corollary 4.5 (*i.i.d.* case and continuous distribution). δ_n is computed from its explicit expression with $K_{4,n} = 9$, $\kappa = 0.99$ (as in the previous numerical applications), and $\lambda_{3,n} = 0.6 \times 9^{3/4} \approx 3.1$. This choice corresponds to a large positive skewness for $|\lambda_{3,n}|$ cannot be larger than $0.621K_{4,n}^{3/4}$ (Pinelis, 2011, Theorem 1). The black line shows the first diagonal $(1 - \Phi(s), 1 - \Phi(s))$.

Figure 4.3 illustrates that, when n is large enough for $\delta_n = O(n^{-1})$ to become negligible compared to the skewness term decreasing in \sqrt{n} , $\lambda_{3,n} \neq 0$ entails a distortion between the true p-value and the approximated one in a one-sided test. As an example, for an approximated p-value of 0.04 and the choices made for $K_{4,n}$, κ , and $\lambda_{3,n}$, the true p-value must be in the interval [0.0397, 0.0429] for n = 5,000 and [0.0404, 0.0407] for n = 25,000.

For a better visualization, Figure 4.3 zooms in the range *pval* between 0.03 and 0.05, which corresponds to test statistics between 1.64 and 1.88, hence above 1. Combined with $\lambda_{3,n} > 0$, it leads to an approximated p-value smaller than the true one. In contrast, Figure 4.4 looks at $\widetilde{pval} \in [0.49, 0.51]$, that is, observed s_n between -0.025 and 0.025. In that range, $|s_n| < 1$ and the positive skewness leads to a positive distortion: the approximated p-value is larger than the true one. The comparison of Figures 4.3 and 4.4, especially for n = 5,000, also exhibits the variation of the magnitude of the shift from the absolute value of $(1 - s^2)\phi(s)$, with the largest gaps for s close to 0.

Finally, we stress that such distortions regarding p-values are specific to one-sided tests. For bilateral or two-sided tests, the skewness of the distribution enters symmetrically in the approximation error and cancels out thanks to the parity of $x \mapsto (1 - x^2)\phi(x)$.²

²In the specific context of this thesis manuscript, it is possible to refer to Equation (5.33) in Chapter 5



Figure 4.4: Same caption as Figure 4.3.



Figure 4.5: Graph of the function $x \mapsto (1 - x^2)\phi(x)$ on the interval [-5, 5].

for the illustration of that simplification.

References

- Abramowitz, M. and Stegun, I. A. (1972), Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables, Vol. 55, National Bureau of Standards, Applied Mathematics Series.
- Adell, J. A. and Lekuona, A. (2008), 'Shortening the distance between Edgeworth and Berry–Esseen in the classical case', *Journal of Statistical Planning and Inference* **138**(4), 1167 1178.
- Bentkus, V. (2003), 'On the dependence of the Berry–Esseen bound on dimension', *Journal* of Statistical Planning and Inference **113**(2), 385 402.
- Bentkus, V. and Götze, F. (1996), 'The Berry–Esseen bound for student's statistic', Ann. Probab. 24(1), 491–503.
- Berry, A. (1941), 'The accuracy of the gaussian approximation to the sum of independent variates', *Transactions of the American Mathematical Society* **49**, 122–136.
- Bhattacharya, R. N. and Ranga Rao, R. (1976), Normal Approximation and Asymptotic Expansions, New York: Wiley.
- Boutsikas, M. V. (2011), 'Asymptotically optimal Berry–Esseen-type bounds for distributions with an absolutely continuous part', *Journal of Statistical Planning and Inference* **141**(3), 1250 – 1268.
- Chernozhukov, V., Chetverikov, D. and Kato, K. (2017), 'Central limit theorems and bootstrap in high dimensions', Ann. Probab. 45(4), 2309–2352.
- Cramer, H. (1962), Random Variables and Probability Distributions, 2 edn, Cambridge University Press.
- de la Peña, V. H., Leung Lai, T. and Shao, Q.-M. (2009), *Self-Normalized Processes: Limit Theory and Statistical Applications*, Probability and Its Aplications, Springer-Verlag, Berlin Heidelberg.
- Efron, B. (1979), 'Bootstrap methods: Another look at the jackknife', Ann. Statist. 7(1), 1–26.
- Esseen, C.-G. (1942), 'On the Liapunoff limit of error in the theory of probability', Arkiv för Matematik, Astronomi och Fysik. .
- Esseen, C.-G. (1945), 'Fourier analysis of distribution functions. a mathematical study of the Laplace-Gaussian law', *Acta Math.* 77, 1–125.
- Esseen, C. G. (1956), 'A moment inequality with an application to the central limit theorem', *Scandinavian Actuarial Journal* **1956**(2), 160–170.
- Gil-Pelaez, J. (1951), 'Note on the inversion theorem', *Biometrika* **38**(3-4), 481–482.
- Goulet, V. (2016), expint: Exponential Integral and Incomplete Gamma Function. R package. URL: https://cran.r-project.org/package=expint

- Hall, P. (1987), 'Edgeworth expansion for Student's t statistic under minimal moment conditions', Ann. Probab. 15(3), 920–931.
- Hall, P. (1992), *The bootstrap and Edgeworth expansion*, Springer series in statistics, Springer-Verlag.
- Kosorok, M. (2006), Introduction to Empirical Processes and Semiparametric Inference, Springer Verlag New York.
- Lahiri, S. N. (2003), *Resampling methods for dependent data*, Springer Science & Business Media.
- Narasimhan, B., Johnson, S. G., Hahn, T., Bouvier, A. and Kiêu, K. (2020), cubature: Adaptive Multivariate Integration over Hypercubes. R package version 2.0.4.1. URL: https://CRAN.R-project.org/package=cubature
- Pinelis, I. (2011), 'Relations between the first four moments'.
- Pinelis, I. and Molzon, R. (2016), 'Optimal-order bounds on the rate of convergence to normality in the multivariate delta method', *Electronic Journal of Statistics* **10**(1), 1001–1063.
- Prawitz, H. (1972), 'Limits for a distribution, if the characteristic function is given in a finite domain', *Scandinavian Actuarial Journal* **1972**(2), 138–154.
- Prawitz, H. (1975), 'On the remainder in the central limit theorem', Scandinavian Actuarial Journal 1975(3), 145–156.
- Raič, M. (2018), 'A multivariate central limit theorem for Lipschitz and smooth test functions', arXiv preprint arXiv:1812.08268.
- Raič, M. (2019), 'A multivariate Berry–Esseen theorem with explicit constants', Bernoulli 25(4A), 2824–2853.
- Senatov, V. V. (2011), 'On the real accuracy of approximation in the central limit theorem', Siberian Mathematical Journal 52(4), 19 38.
- Shevtsova, I. (2010), 'Refinement of estimates for the rate of convergence in Lyapunov's theorem', *Dokl. Akad. Nauk* **435**(1), 26–28.
- Shevtsova, I. (2012), 'Moment-type estimates with asymptotically optimal structure for the accuracy of the normal approximation', Annales Mathematicae et Informaticae 39, 241–307.
- Shevtsova, I. G. (2013), 'On the absolute constants in the berry–esseen inequality and its structural and nonuniform improvements', *Informatika i Ee Primeneniya [Informatics and its Applications]* 7(1), 124–125.
- Ushakov, N. G. (2011), Selected Topics in Characteristic Functions, De Gruyter, Berlin, Boston.
- Ushakov, N. G. and Ushakov, V. G. (1999), 'Some inequalities for characteristic functions of densities with bounded variation', *Preprint series. Statistical Research Report http://urn.* nb. no/URN: NBN: no-23420.

van der Vaart, A. (2000), Asymptotics Statistics, Cambridge University Press.

- van der Vaart, A. and Wellner, J. (1996), Weak Convergence of Empirical Processes: with Applications to Statistics, Springer-Verlag New York.
- Zhilova, M. (2020), 'New Edgeworth-type expansions with finite sample guarantees', arXiv preprint arXiv:2006.03959.

Appendix 4.A Proof of the main results

In this section, we prove Theorems 4.1 and 4.2. All the proofs start with a so-called "smoothing inequality", which we present and prove in Subsection 4.A.1. Based on this inequality, Subsection 4.A.2 presents the common structure of all the proofs. The remaining subsections are devoted to proving the theorems themselves.

4.A.1 A smoothing inequality

The result given in Lemma 4.1 helps control the distance between the cumulative distribution function $\mathbb{P}(S_n \leq x)$ and $G_n(x) := \Phi(x) + \frac{\lambda_{3,n}}{6\sqrt{n}}(1-x^2)\varphi(x)$ in terms of their respective Fourier transforms.

Lemma 4.1. For every $t_0 \in (0, 1]$ and every T > 0, we have

$$\sup_{x \in \mathbb{R}} \left| \mathbb{P} \left(S_n \le x \right) - G_n(x) \right| \le \Omega_1(t_0, T, |\lambda_{3,n}| / \sqrt{n}) + \Omega_2(t_0, T) + \Omega_3(t_0, T, \lambda_{3,n} / \sqrt{n}), \quad (4.8)$$

where

$$\begin{split} \Omega_1(t_0, T, v) &:= 2 \int_0^{t_0} \left| \Psi(t) - \frac{i}{2\pi t} \right| e^{-(Tt)^2/2} \left(1 + \frac{v|Tt|^3}{6} \right) dt \\ &\quad + \frac{1}{\pi} \int_{t_0}^{+\infty} \frac{e^{-(Tt)^2/2}}{t} \left(1 + \frac{v|Tt|^3}{6} \right) dt, \\ \Omega_2(t_0, T) &:= 2 \int_{t_0}^{1} |\Psi(t)| \left| f_{S_n}(Tt) \right| dt, \\ \Omega_3(t_0, T, v) &:= 2 \int_0^{t_0} |\Psi(t)| \left| f_{S_n}(Tt) - e^{-(Tt)^2/2} \left(1 - \frac{vi(Tt)^3}{6} \right) \right| dt, \end{split}$$

and $\Psi(t) := \frac{1}{2} \left(1 - |t| + i \left[(1 - |t|) \cot(\pi t) + \frac{\operatorname{sign}(t)}{\pi} \right] \right) \mathbb{1} \{ |t| \le 1 \}.$

In the following, we resort to Equations (I.29) and (I.30) of Prawitz (1975) which claim that the function Ψ satisfies

$$\left|\Psi(t)\right| \le \frac{1.0253}{2\pi|t|} \text{ and } \left|\Psi(t) - \frac{i}{2\pi t}\right| \le \frac{1}{2}\left(1 - |t| + \frac{\pi^2}{18}t^2\right).$$
 (4.9)

Proof of Lemma 4.1. Let us denote by "p.v. \int " Cauchy's principal value, defined by

p.v.
$$\int_{-a}^{a} f(u) du := \lim_{x \to 0, x > 0} \int_{-a}^{-x} f(u) du + \int_{x}^{a} f(u) du$$
,

where f is a measurable function on $[-a, a] \setminus \{0\}$ for a given a > 0. In the following, we use the following inequalities, which are due to Prawitz (1972), where F is the cumulative distribution function of S_n and $f = f_{S_n}$ its characteristic function,

$$\lim_{y \to x, y > x} F(y) \leq \frac{1}{2} + \text{p.v.} \int_{-U}^{+U} e^{-ixu} \frac{1}{U} \Psi\left(\frac{u}{U}\right) f(u) du,$$
$$\lim_{y \to x, y < x} F(y) \geq \frac{1}{2} + \text{p.v.} \int_{-U}^{+U} e^{-ixu} \frac{1}{U} \Psi\left(\frac{-u}{U}\right) f(u) du.$$

Therefore,

$$F(x) - G_n(x) \le \frac{1}{2} + \text{p.v.} \int_{-U}^{+U} e^{-ixu} \frac{1}{U} \Psi\left(\frac{u}{U}\right) f(u) du - G_n(x)$$
(4.10)

$$F(x) - G_n(x) \ge \frac{1}{2} + \text{p.v.} \int_{-U}^{+U} e^{-ixu} \frac{1}{U} \Psi\left(\frac{-u}{U}\right) f(u) du - G_n(x).$$
(4.11)

Note that the Gil-Pelaez inversion formula (see Gil-Pelaez (1951)) is valid for any boundedvariation function. Formally, for every bounded-variation function $G(x) = \int_{-\infty}^{x} g(t)dt$, denoting the Fourier transform of a given function g by $\check{g} := \int_{-\infty}^{+\infty} e^{ixu}g(u)du$, we have

$$G(x) = \frac{1}{2} + \frac{i}{2\pi} \text{ p.v.} \int_{-\infty}^{+\infty} e^{-ixu} \check{g}(u) du.$$
(4.12)

Therefore, applying Equation (4.12) to the function $G_n(x) := \Phi(x) + \lambda_{3,n}(1-x^2)\varphi(x)/(6\sqrt{n})$ whose (generalized) density has the Fourier transform $(1 - \lambda_{3,n}ix^3/(6\sqrt{n}))e^{-x^2/2}$, we get

$$G_n(x) = \frac{1}{2} + \frac{i}{2\pi} \text{ p.v.} \int_{-\infty}^{+\infty} e^{-ixu} \left(1 - \frac{\lambda_{3,n}}{6\sqrt{n}} iu^3\right) e^{-u^2/2} \frac{du}{u}.$$

Combining this equality with the bounds (4.10) and (4.11) and using the triangular inequality, we get the claimed result (4.8).

4.A.2 Outline of the proofs of Theorems 4.1 and 4.2

The following subsections show more general formulations of Theorems 4.1 and 4.2. They yield the two theorems as stated in the body of the article for a particular choice of a tuning parameter ε . Besides, Theorems 4.1 and 4.2 are stated under two possible sets of assumptions: Assumption 4.1 (*i.n.i.d.*) or Assumption 4.2 (*i.i.d.*). The proof is done separately for each theorem and each assumption, resulting in four different proofs. However, they share the same structure. This preamble aims at presenting that structure and pointing out the similarities and differences between the four proofs.

Structure

There are four distinct settings and associated theorems:

- 1. moment conditions only and *i.n.i.d.* observations ("no-continuity, *i.n.i.d.*"), in this setting, the bound δ_n on $\Delta_{n,E}$ is proved in Theorem 4.3 (Section 4.A.3);
- 2. moment conditions only and *i.i.d.* observations ("no-continuity, *i.i.d.*"), in this setting, the bound δ_n on $\Delta_{n,E}$ is proved in Theorem 4.4 (Section 4.A.4);
- 3. additional regularity conditions and *i.n.i.d.* observations ("continuity, *i.n.i.d.*"),³ in this setting, the bound δ_n on $\Delta_{n,E}$ is proved in Theorem 4.5 (Section 4.A.5);
- 4. additional regularity conditions and *i.i.d.* observations ("continuity, *i.i.d.*"), in this setting, the bound δ_n on $\Delta_{n,E}$ is proved in Theorem 4.6 (Section 4.A.6).

³The regularity conditions control the tail behavior of the characteristic function f_{S_n} of the standardized sum S_n in the *i.n.i.d.* case (resp. of the characteristic function of one standardized observation X_n/σ_n in the *i.i.d.* case). Although the formal link with an absolutely continuous distribution (or a component thereof) with respect to Lebesgue's measure concerns the *i.i.d.* case, we use the term "continuity" setting to refer to those additional assumptions both in *i.i.d.* and *i.n.i.d.* cases.

Note that Theorem 4.5 and Theorem 4.6 do not require the additional regularity assumptions but hold under moment conditions *only* (Assumption 4.1 for the former and Assumption 4.2 for the latter). Together, they yield Theorem 4.2 presented in the body of the article. Nonetheless, as explained in the discussion of Theorem 4.2, the leading term of the obtained bound δ_n depends on the tail behavior of f_{S_n} and will result in an improvement provided additional regularity conditions on f_{S_n} (see Corollary 4.3 in the *i.n.i.d.* case and Corollary 4.5 in the *i.i.d.* case). That is why we refer nonetheless to the settings of Theorem 4.5 and Theorem 4.6 as "continuity" settings.

Furthermore, in each of those four settings, some terms in the bounds can be improved under the assumption $\lambda_{3,n} = 0$ (unskewed case). We do not detail the related changes here (see the following subsections) to focus on the main structure.

The four proofs start by applying Lemma 4.1 and, then, deal separately with the three terms of Equation (4.8): $\Omega_1(t_0, T, |\lambda_{3,n}|/\sqrt{n}), \Omega_2(t_0, T), \text{ and } \Omega_3(t_0, T, \lambda_{3,n}/\sqrt{n}).$

 $\lambda_{3,n}$ depends on the distribution of the $(X_i)_{i=1,\dots,n}$. In contrast, $t_0 \in (0, 1]$ and T > 0 are free parameters. In the four settings, t_0 remains unchanged and set at $1/\pi \approx 0.32$. However, the choice for T differs between the "no-continuity" (1. and 2.) and "continuity" (3. and 4.) settings. As explained below, this is precisely this change that enables to modify the leading term in the bound δ_n and obtain a leading term that depends on the tail behavior of f_{S_n} in Theorems 4.5 and 4.6.

First term Ω_1

The upper bound on Ω_1 is formally the same in the four settings (as a function of an arbitrary T; its actual value depends on the choice of T). Lemma 4.4 in Section 4.B.1 shows that, for every T > 0,

$$\Omega_1(1/\pi, T, |\lambda_{3,n}|/\sqrt{n}) \le \frac{1.2533}{T} + \frac{0.3334|\lambda_{3,n}|}{T\sqrt{n}} + \frac{14.1961}{T^4} + \frac{4.3394|\lambda_{3,n}|}{T^3\sqrt{n}}$$

In the "no-continuity" setting, T is chosen of the order of \sqrt{n} :

$$T = \frac{2\pi\sqrt{n}}{\widetilde{K}_{3,n}} \quad (\text{``no-continuity'' choice of } T).$$

Consequently, the first term in Ω_1 will enter in the dominant term (decreasing at rate \sqrt{n}) of the bound δ_n on $\Delta_{n,E}$ in the settings 1. and 2. On the contrary, in the "continuity" setting, we choose a larger T (the "no-continuity" one at the power 4), of the order of n^2 :

$$T = \frac{16\pi^4 n^2}{\widetilde{K}_{3,n}^4} \quad (\text{"continuity" choice of } T).$$

That choice enables to "kill" the term Ω_1 in the sense that, under proper regularity assumptions on f_{S_n} , the bound on $\Omega_1(1/\pi, T, |\lambda_{3,n}|/\sqrt{n})$ becomes negligible in δ_n (the dominant term 1.2533/T decreases to 0 at the speed n^2 for the "continuity" choice of T).

Second term Ω_2

The treatment of the second term Ω_2 coming from the smoothing inequality differs between the "no-continuity" (1. and 2.) and "continuity" (3. and 4.) settings.

In "no-continuity" settings 1. and 2. The bound on Ω_2 derives from Lemma 4.2 that enables to control the modulus of the characteristic function f_{S_n} of S_n in the relevant interval $[1/\pi, 1]$. We obtain (see Equation (4.15) below)

$$\Omega_2(1/\pi, T) \le \frac{1.2187}{T^2} + \frac{67.0415}{T^4}.$$

For the choice of T in the "no-continuity" setting, the term Ω_2 is thus negligible compared to the dominant term in the bound δ_n , which decreases at rate \sqrt{n} .

In "continuity" settings 3. and 4. Recall that $\Omega_2(1/\pi, T) = 2 \int_{\pi^{-1}}^{1} |\Psi(t)| |f_{S_n}(Tt)| dt$. Using the upper bound on the modulus of the kernel function Ψ (Equation (4.9), Prawitz (1975)), the term Ω_2 can be bounded by an integral whose only unknown term is the modulus of f_{S_n} . This is where we recover the interesting integral term involving f_{S_n} of Theorem 4.2. In other words, in "continuity" settings 3. and 4., we do *not* want to apply Lemma 4.2 used in the "no-continuity" settings to control Ω_2 , but we keep its expression involving the integral of $|f_{S_n}|$ in order to use the additional regularity assumptions on the characteristic function of S_n . Indeed, a change of variables and the first part of Equation (4.9) give

$$\Omega_2(1/\pi, T) \le \frac{1.0253}{\pi} \int_{T/\pi}^T \frac{|f_{S_n}(u)|}{u} du$$

From there, the additional regularity assumptions on f_{S_n} provide an explicit control on Ω_2 .

Third term Ω_3

The control of the third term Ω_3 is more involved compared to Ω_1 and Ω_2 . It is done by controlling separately several terms whose sum provides an upper bound on Ω_3 (see the next four subsections for details and the definitions of those terms): $\Omega_3(1/\pi, T, \lambda_{3,n}/\sqrt{n}) \leq$

$$I_{3,1}(T) + I_{3,2}(T) + I_{3,3}(T)$$
(setting 1.)

$$I_{4,1}(T) + I_{4,2}(T) + I_{4,3}(T)$$
(setting 2.)

$$I_{5,1}(T) + I_{5,2}(T) + I_{5,3}(T) + I_{5,4}(T) + I_{5,5}(T)$$
(settings 3. and 4.)

The upper bounds are distinct and derived in each of the four settings although several of the terms are controlled similarly and some are even equal (e.g., $I_{5,1}(T) = I_{4,1}(T)$).⁴

4.A.3 Proof of Theorem 4.1 under Assumption 4.1

In this section, we state and prove a more general theorem (Theorem 4.4 below). We recover Theorem 4.1 when we set $\varepsilon = 0.1$ (numerical computations of $e_1(0.1)$ give the upper bound $e_1(0.1) \leq 1.012$).

 $^{^{4}}$ We acknowledge that, for future work, some modifications and harmonization in the notation of those terms might ease the readability of the proofs. We hope this preamble subsection goes in that direction.

Theorem 4.3 (One-term Edgeworth expansion under Assumption 4.1). Under Assumption 4.1, for every $\varepsilon \in (0, 1/3)$ and every $n \ge 1$, we have the bound

$$\Delta_{n,E} \leq \frac{0.1995 \,\widetilde{K}_{3,n}}{\sqrt{n}} + \frac{1}{n} \Biggl\{ 0.031 \,\widetilde{K}_{3,n}^2 + 0.327 \,K_{4,n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) + 0.054 \,|\lambda_{3,n}|\widetilde{K}_{3,n} + 0.037 \,e_1(\varepsilon)|\lambda_{3,n}|^2 \Biggr\} + r_{1,n},$$

$$(4.13)$$

where $e_1(\varepsilon)$ is given in Equation (4.42) and $r_{1,n}$ is given in Equation (4.22). Note that $r_{1,n} = O(n^{-5/4})$ as soon as $K_{4,n} = O(1)$. If $\mathbb{E}[X_i^3] = 0$ for every i = 1, ..., n, the upper bound reduces to

$$\frac{0.1995\widetilde{K}_{3,n}}{\sqrt{n}} + \frac{1}{n} \left\{ 0.031\widetilde{K}_{3,n}^2 + 0.327K_{4,n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) \right\} + r_{1,n},$$
(4.14)

with $r_{1,n} = O(n^{-3/2})$ when $K_{4,n} = O(1)$.

We follow the proof strategy initiated by Prawitz (1975) and complemented among others by Shevtsova (2012). The method of proof starts from Lemma 4.1. Note that in Prawitz (1975) or Shevtsova (2012), Lemma 4.1 is used with $\Phi(x)$ instead of $G_n(x)$ (that is, these authors are only interested in the canonical Berry-Esseen inequality). It is shown in Section 4.B.1 how to upper bound the term $\Omega_1(t_0, T, |\lambda_{3,n}|/\sqrt{n})$. There remains to control $\Omega_2(t_0, T)$ and $\Omega_3(t_0, T, \lambda_{3,n}/\sqrt{n})$. The former is handled as in Shevtsova (2012) while the latter is tackled differently.

As in the classical Berry-Esseen setup, T is chosen of the order of \sqrt{n} . The parameter t_0 is chosen as in Prawitz (1975): a range of values is admissible, in particular t_0 must satisfy $1/(2\pi) \leq t_0 \leq 1/(6\pi\chi_1)$ where $\chi_1 := \sup_{x>0} x^{-3} |\cos(x) - 1 + x^2/2| \approx 0.099162$. Given the numerical evidence in Prawitz (1975), the choice $t_0 = 1/\pi$ is reasonable.

Let $\xi_n := K_{3,n}/\sqrt{n}$. Based on the upper bound on Ω_1 , we would like to pick T as large as possible. However, Lemma 4.2 below restricts the range of informative Ts to control Ω_2 . More specifically, we show in Section 4.A.3 that a suitable choice for T is

$$T = \frac{2\pi}{\xi_n} = \frac{2\pi}{\widetilde{K}_{3,n}/\sqrt{n}} = \frac{2\pi\sqrt{n}}{\widetilde{K}_{3,n}}$$

Bound on Ω_2

In this section, we control $\Omega_2(1/\pi, T) = 2 \int_{1/\pi}^1 |\Psi(t)| |f_{S_n}(Tt)| dt$. To reach this goal, we use the following lemma which is a consequence of Lemma 2 in Shevtsova (2010) or Theorem 2.2 in Shevtsova (2012) (with $\delta = 1$):

Lemma 4.2. Let $t_1^* = \theta_1^*/(2\pi)$ where θ_1^* is the unique root in $(0, 2\pi)$ of the equation $\theta^2 + 2\theta \sin(\theta) + 6(\cos(\theta) - 1) = 0.^5$ For T such that $t_1^*/\xi_n \leq Tt \leq 2\pi/\xi_n$ for every $t \in [t_1^*, 1]$, we have

$$|f_{S_n}(Tt)| \le e^{-(Tt)^2/2 + \chi_1 \xi_n |Tt|^3} \text{ for } t \in [1/\pi, t_1^*]$$

$$|f_{S_n}(Tt)| \le e^{-(1 - \cos(\xi_n Tt))/\xi_n^2} \text{ for } t \in [t_1^*, 1].$$

 ${}^{5}t_{1}^{*} \approx 0.64.$

Therefore,

$$\int_{1/\pi}^{1} |\Psi(t)| |f_{S_n}(Tt)| dt \le \int_{1/\pi}^{t_1^*} |\Psi(t)| e^{-(Tt)^2/2 + \chi_1 \xi_n |Tt|^3} dt + \int_{t_1^*}^{1} |\Psi(t)| e^{-(1 - \cos(\xi_n Tt))/\xi_n^2} dt$$

Proof of Lemma 4.2: Applying Theorem 2.2 in Shevtsova (2012) with $\delta = 1$, we get for all $u \in \mathbb{R}$

$$|f_{S_n}(u)| \le \exp\left(-\psi(u,\epsilon_n)\right)$$

where $\epsilon_n := n^{-1/2} \widetilde{K}_{3,n}$, and, for any real $u, \epsilon > 0$

$$\psi(u,\epsilon) := \begin{cases} t^2/2 - \chi_1 \epsilon |t|^3, & \text{for } |t| < \theta_1^* \epsilon^{-1}, \\ \frac{1 - \cos(\epsilon t))}{\epsilon^2}, & \text{for } \theta_1^* \epsilon^{-1} \le |t| \le 2\pi \epsilon^{-1}, \\ 0, & \text{for } |t| > 2\pi \epsilon^{-1}. \end{cases}$$

Multiplying by $|\Psi|$, integrating from $1/\pi$ to 1 and separating the two cases yields the claimed inequality. \Box

As already mentioned, we want to take T as large as possible to control Ω_1 . However, the proof of Lemma 4.2 implies that choosing $T > 2\pi/\xi_n$ would resort in a useless bound on Ω_2 . This is why we pick $T = 2\pi/\xi_n$. We can thus write $\chi_1\xi_n|Tt|^3 = 2\pi\chi_1(Tt)^2|t|$ for $t \in [1/\pi, t_1^*]$ and $\cos(\xi_n Tt) = \cos(2\pi t)$ for $t \in [t_1^*, 1]$. Combining this with Lemma 4.2 yields

$$\begin{split} \int_{1/\pi}^{1} |\Psi(t)| \, |f_{S_n}(Tt)| dt &\leq \int_{1/\pi}^{t_1^*} |\Psi(t)| e^{-\frac{(Tt)^2}{2}(1-4\pi\chi_1|t|)} dt + \int_{t_1^*}^{1} |\Psi(t)| e^{-T^2(1-\cos(2\pi t))/(4\pi^2)} dt \\ &= \frac{I_{2,1}(T)}{2T^4} + \frac{I_{2,2}(T)}{2T^2}, \end{split}$$

where

$$I_{2,1}(T) := T^4 \int_{1/\pi}^{t_1^*} 2|\Psi(t)| e^{-\frac{(Tt)^2}{2}(1-4\pi\chi_1|t|)} dt,$$

$$I_{2,2}(T) := T^2 \int_{t_1^*}^{1} 2|\Psi(t)| e^{-T^2(1-\cos(2\pi t))/(4\pi^2)} dt.$$

Note that the difference in the two exponents of T in the above definitions may seem surprising as these two integrals look similar. However they have very different behaviors since the first one decays much faster than the second one. In line with Section 4.B.1, we compute numerically these integrals using the R package cubature (Narasimhan et al., 2020) and optimize them using the optimize function with the L-BFGS-B method.⁶ This gives

$$\sup_{T \ge 0} I_{2,1}(T) \le 67.0415, \text{ and } \sup_{T \ge 0} I_{2,2}(T) \le 1.2187.$$

Finally, we arrive at

$$\Omega_2(1/\pi, T) = 2 \int_{1/\pi}^1 |\Psi(t)| \, |f_{S_n}(Tt)| dt \le \frac{67.0415}{T^4} + \frac{1.2187}{T^2}. \tag{4.15}$$

 $^{^{6}\}mathrm{We}$ would be happy to share the code upon request.

Bound on Ω_3

In this section, we bound the third term of Equation (4.8), which is

$$\Omega_3(t_0, T, \lambda_{3,n}/\sqrt{n}) = 2 \int_0^{t_0} |\Psi(t)| \left| f_{S_n}(Tt) - e^{-(Tt)^2/2} \left(1 - \frac{\lambda_{3,n} i(Tt)^3}{6\sqrt{n}} \right) \right| dt.$$

Whenever $t_0T = T/\pi \ge \sqrt{2\varepsilon}(n/K_{4,n})^{1/4}$ – or equivalently when $n \ge \varepsilon^2 K_{3,n}^4/(4K_{4,n})$ – we can write

$$\Omega_3(t_0, T, \lambda_{3,n}/\sqrt{n}) \le I_{3,1}(T) + I_{3,2}(T) + I_{3,3}(T),$$

where

$$\begin{split} I_{3,1}(T) &:= \frac{2}{T} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \left(1 - \frac{iu^3 \lambda_{3,n}}{6\sqrt{n}} \right) \right| du, \\ I_{3,2}(T) &:= \frac{2}{T} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{t_0 T} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \right| du, \\ I_{3,3}(T) &:= \frac{2}{T} \frac{|\lambda_{3,n}|}{6\sqrt{n}} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{t_0 T} |\Psi(u/T)| e^{-u^2/2} |u|^3 du. \end{split}$$

Note that when $n < \varepsilon^2 K_{3,n}^4/(4K_{4,n})$, we have the better inequality $\Omega_3(t_0, T, \lambda_{3,n}/\sqrt{n}) < I_{3,1}(T)$, so that our reasoning is still valid even in this case as our bounds are all positive.

The integrand of $I_{3,1}(T)$ can be controlled with the help of Lemma 4.5, which enables us to write

$$I_{3,1}(T) \leq \frac{K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) J_1\left(4, 0, \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}, T\right) + \frac{e_{1,n}(\varepsilon)}{36} \frac{|\lambda_{3,n}|^2}{n} J_1\left(6, 0, \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}, T\right) + I_{3,1,3}$$
(4.16)

where

$$I_{3,1,3}(T) := \frac{2}{T} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} |\Psi(u/T)| e^{-u^2/2} R_{1,n}(u,\varepsilon) du,$$

and J_1 is defined in Equation (4.55). Using Equation (4.55), we obtain the bounds $J_1(4, 0, +\infty, T) \leq 0.327$ and $J_1(6, 0, +\infty, T) \leq 1.306$. Besides by the first inequality in (4.9), we get

$$I_{3,1,3}(T) \le \frac{1.0253}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} R_{1,n}(u,\varepsilon) du.$$

We finally get from Equation (4.16)

$$I_{3,1}(T) \leq \frac{0.327K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) + \frac{1.306e_{1,n}(\varepsilon)}{36} \frac{|\lambda_{3,n}|^2}{n} + \frac{1.0253}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} R_{1,n}(u,\varepsilon) du.$$
(4.17)

We now handle $I_{3,3}(T)$. We remark that the following set of inequalities is valid for p = 3

$$\begin{split} I_{3,3}(T) &= \frac{|\lambda_{3,n}|}{3\sqrt{n}} J_1\Big(3, \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}, t_0 T, T\Big) \\ &\leq \frac{1.0253|\lambda_{3,n}|}{6\pi\sqrt{n}} \int_{\sqrt{2\varepsilon} (n/K_{4,n})^{1/4}}^{2\sqrt{n}/\widetilde{K}_{3,n}} u^2 e^{-u^2/2} du \\ &= \frac{1.0253|\lambda_{3,n}|}{6\pi\sqrt{n}} \int_{\varepsilon (n/K_{4,n})^{1/2}}^{2n/\widetilde{K}_{3,n}^2} (2v)^{(2-1)/2} e^{-v} \frac{dv}{\sqrt{2v}} \\ &= \frac{1.0253|\lambda_{3,n}|}{6\pi\sqrt{n}} \Big(\Gamma\Big(2/2, \varepsilon (n/K_{4,n})^{1/2}\Big) - \Gamma\Big(2/2, 2n/\widetilde{K}_{3,n}^2\Big) \Big) \\ &= O\Big(\Gamma\Big(p/2, \varepsilon (n/K_{4,n})^{1/2}\Big) - \Gamma\Big(p/2, 2n/\widetilde{K}_{3,n}^2\Big) \Big) \\ &= O\Big(n^{p/4 - 1/2} e^{-\varepsilon\sqrt{n}/\sqrt{K_{4,n}}} \Big), \end{split}$$

where we apply the change of variable $v = u^2/2$, and take advantage of the asymptotic expansion $\Gamma(a, x) = x^{a-1}e^{-x}(1 + O((a-1)/x))$ which is valid for every fixed a in the regime $x \to \infty$, see Equation (6.5.32) in Abramowitz and Stegun (1972).

In the following section, we show that $I_{3,2}(T)$ decays exponentially with n too, so that we obtain

$$\Omega_{3}(t_{0}, T, \lambda_{3,n}/\sqrt{n}) \leq I_{3,1}(T) + I_{3,2}(T) + I_{3,3}(T)$$

$$\leq \frac{0.327K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^{2}}\right) + \frac{1.306e_{1,n}(\varepsilon)}{36} \frac{|\lambda_{3,n}|^{2}}{n} + O(n^{-5/4}).$$
(4.18)

Bound on $I_{3,2}(T)$.

Let t be a real in the interval $[\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, t_0T]$. As in the proof of Lemma 2.7 in Shevtsova (2012) with $\delta = 1$, using the fact that for every $i = 1, \ldots, n$

$$\max\left\{|f_{P_{X_i}}(t)|, \exp\left(-\frac{t^2\sigma_i^2}{2}\right)\right\} \le \exp\left(-\frac{t^2\sigma_i^2}{2} + \frac{\chi_1 t^3(\mathbb{E}[|X_i|^3] + \mathbb{E}[|X_i|]\sigma_i^2)}{B_n^3}\right),$$

we get

$$\begin{aligned} \left| f_{S_n}(t) - e^{-t^2/2} \right| &\leq \sum_{i=1}^n \left| f_{P_{X_i}} \left(\frac{t}{B_n} \right) - e^{-\frac{t^2 \sigma_i^2}{2B_n^2}} \right| e^{\frac{t^2 \sigma_i^2}{2B_n^2}} e^{-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \sum_{l=1}^n \left(\mathbb{E}[|X_l|^3] + \mathbb{E}[|X_i|] \sigma_i^2 \right)}{B_n^3} \\ &= \sum_{i=1}^n \left| f_{P_{X_i}} \left(\frac{t}{B_n} \right) - e^{-\frac{t^2 \sigma_i^2}{2B_n^2}} \right| e^{-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2 \sigma_i^2}{2B_n^2}}. \end{aligned}$$

By Equation (4.44), we have $\max_{1 \le i \le n} \sigma_i^2 \le B_n^2 \times (K_{4,n}/n)^{1/2}$ so that we obtain

$$\left| f_{S_n}(t) - e^{-t^2/2} \right| \le \sum_{i=1}^n \left| f_{P_{X_i}}\left(\frac{t}{B_n}\right) - e^{-\frac{t^2 \sigma_i^2}{2B_n^2}} \right| e^{-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2}{2} \sqrt{\frac{K_{4,n}}{n}}}.$$

Applying Lemma 2.8 in Shevtsova (2012), we get that for every variable X such that $\mathbb{E}[|X|^3]$ is finite, $|f(t) - e^{-\sigma^2 t^2}| \leq \mathbb{E}[|X|^3] \times |t|^3/6$. Therefore,

$$\left| f_{S_n}(t) - e^{-t^2/2} \right| \leq \sum_{i=1}^n \frac{\mathbb{E}[|X_i|^3]}{6B_n^3} |t|^3 \exp\left(-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2}{2} \sqrt{\frac{K_{4,n}}{n}}\right)$$
$$= \frac{K_{3,n}}{6\sqrt{n}} |t|^3 \exp\left(-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2}{2} \sqrt{\frac{K_{4,n}}{n}}\right).$$
(4.19)

Recalling that $t_0 = 1/\pi$, and integrating the latter equation, we have

$$I_{3,2}(T) = \frac{2}{T} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{t_0 T} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \right| du,$$

$$\leq \frac{K_{3,n}}{3\sqrt{nT}} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T/\pi} |\Psi(u/T)| u^3 \exp\left(-\frac{u^2}{2} + \frac{\chi_1 |u|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{u^2}{2} \sqrt{\frac{K_{4,n}}{n}}\right) du$$

$$= \frac{K_{3,n}}{3\sqrt{n}} J_2 \Big(3, \sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, T/\pi, \widetilde{K}_{3,n}, K_{4,n}, T, n\Big)$$

$$\leq \frac{1.0253 K_{3,n}}{6\pi \sqrt{n}} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T/\pi} u^2 \exp\left(-\frac{u^2}{2} + \frac{\chi_1 |u|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{u^2}{2} \sqrt{\frac{K_{4,n}}{n}}\right) du$$
(4.20)

where J_2 is as defined in Equation (4.54).

Improved bound on $I_{3,2}(T)$ under the assumption that $\mathbb{E}[X_i^3] = 0$ for all i

When $\mathbb{E}[X_i^3] = 0$ for all i = 1, ..., n, the bound on $I_{3,2}(T)$ can be further improved. The proof mostly follows the reasoning of Section 4.A.3, with suitable modifications.

First, using a Taylor expansion of order 3 of $f_{P_{X_i}}$ around 0 (with explicit Lagrange remainder) and the inequality $|e^{-x} - 1 + x| \leq x^2/2$, we can claim for every real t

$$\left| f_{P_{X_i}}(t) - e^{-t^2 \sigma_i^2 / 2} \right| \le \frac{t^4 \gamma_i}{24} + \frac{\sigma_i^4 t^4}{8} \le \frac{t^4 \gamma_i}{6}.$$

Reasoning as in the proof of Lemma 2.7 in Shevtsova (2012) with $\delta = 1$, we obtain

$$\left| f_{S_n}(t) - e^{-t^2/2} \right| \leq \sum_{i=1}^n \frac{t^4 \gamma_i}{6B_n^4} \exp\left(-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2}{2} \sqrt{\frac{K_{4,n}}{n}}\right)$$
$$\leq \frac{K_{4,n}}{6n} t^4 \exp\left(-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2}{2} \sqrt{\frac{K_{4,n}}{n}}\right).$$

Plugging this into the definition of $I_{3,2}(T)$, we can write

$$\begin{split} I_{3,2}(T) &= \frac{2}{T} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{t_0 T} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \right| du, \\ &\leq \frac{K_{4,n}}{3nT} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T/\pi} |\Psi(u/T)| \, u^4 \exp\left(-\frac{u^2}{2} + \frac{\chi_1 |u|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{u^2}{2} \sqrt{\frac{K_{4,n}}{n}}\right) du \\ &\leq \frac{K_{4,n}}{3n} J_2 \Big(4, \sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, T/\pi, \widetilde{K}_{3,n}, K_{4,n}, T, n), \\ &\leq \frac{1.0253K_{4,n}}{6\pi n} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T/\pi} u^3 \exp\left(-\frac{u^2}{2} + \frac{\chi_1 |u|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{u^2}{2} \sqrt{\frac{K_{4,n}}{n}}\right) du. \end{split}$$

Conclusion: end of the proof of Theorem 4.3.

We start from Equation (4.8). The quantity Ω_1 is bounded by Lemma 4.4, the quantity Ω_2 in Equation (4.15) and the quantity Ω_3 in Equation (4.18). As a result,

$$\begin{split} \sup_{x \in \mathbb{R}} |\mathbb{P} \left(S_n \leq x \right) - G_n(x) | &\leq \Omega_1(t_0, T, |\lambda_{3,n}| / \sqrt{n}) + \Omega_2(t_0, T) + \Omega_3(t_0, T, \lambda_{3,n} / \sqrt{n}) \\ &\leq \frac{1.2533}{T} + \frac{0.3334 |\lambda_{3,n}|}{T\sqrt{n}} + \frac{14.1961}{T^4} + \frac{4.3394 |\lambda_{3,n}|}{T^3\sqrt{n}} \\ &\quad + \frac{67.0415}{T^4} + \frac{1.2187}{T^2} + I_{3,1}(T) + I_{3,2}(T) + I_{3,3}(T) \\ &\leq \frac{1.2533\widetilde{K}_{3,n}}{2\pi\sqrt{n}} + \frac{0.3334 |\lambda_{3,n}| \widetilde{K}_{3,n}}{2\pi n} + \frac{1.2187\widetilde{K}_{3,n}^2}{4\pi n} + \frac{0.327K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) \\ &\quad + \frac{1.306e_{1,n}(\varepsilon) |\lambda_{3,n}|^2}{36n} + r_{1,n}, \end{split}$$

where

1

$$r_{1,n} := \frac{(14.1961 + 67.0415)\widetilde{K}_{3,n}^4}{16\pi^4 n^2} + \frac{4.3394|\lambda_{3,n}|\widetilde{K}_{3,n}^3}{8\pi^3 n^2} + \frac{K_{3,n}}{3\sqrt{n}} J_2 \Big(3, \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}, T/\pi, \widetilde{K}_{3,n}, K_{4,n}, T, n\Big) + I_{3,3}(T) + \frac{1.0253}{\pi} \int_0^{\sqrt{2\varepsilon} (n/K_{4,n})^{1/4}} u e^{-u^2/2} R_{1,n}(u,\varepsilon) du.$$
(4.22)

Given the definition of $R_{1,n}(u,\varepsilon)$, the last term in (4.22) can be written as a sum of elements depending on ε and on positive powers of $\frac{K_{4,n}}{n}$ and $\frac{|\lambda_{3,n}|}{\sqrt{n}}$. Since $K_{4,n}$ and $|\lambda_{3,n}|$ are bounded by assumption, we can see based on the definition of $R_{1,n}(u,\varepsilon)$ that the last (and dominant) term in the definition of $r_{1,n}$ comes from $U_{1,2,n}$, defined in Equation (4.43) and gives the rate $O(n^{-5/4})$.

When $\mathbb{E}[X_i^3] = 0$ for every i = 1, ..., n, we have $\lambda_{3,n} = 0$. This removes the associated terms. The second line in (4.22) becomes $\frac{K_{4,n}}{3n}J_2\left(4,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T/\pi,\widetilde{K}_{3,n},K_{4,n},T,n\right)$. The dominant term in $r_{1,n}$ which stems from Equation (4.43) is null whenever $\sum_{i=1}^{n} |\mathbb{E}[X_i^3]| = 0$. Under the new assumption, this term disappears and the next term becomes the dominant one in the remainder $r_{1,n}$. We finally obtain the bound $r_{1,n} = O(n^{-3/2}).$

The remainder term $r_{1,n}$ as defined in Equation (4.22) is fully explicit although the expression involves a number of complicated integrals. In the practical use of our bounds on $\Delta_{n,E}$ and $\Delta_{n,B}$ (see Examples 4.2 and 4.6), we stick to that expression to obtain bounds that are as sharp as possible using numerical integration to compute the different terms. Nonetheless, to make the analysis of the asymptotic behavior of $r_{1,n}$ more transparent, it is possible to further bound the different terms involved in Equation (4.22) by simpler quantities. This is done in the subsection below for $r_{1,n}$ in the setting of Assumption 4.1 (setting 1. "no-continuity *i.n.i.d.*") and with possible skewness (i.e., $\lambda_{3,n} \neq 0$ a priori). Similar analyses could be done to make more transparent the asymptotic behavior of the remainder terms $r_{1,n}$ and $r_{2,n}$ under the different settings.

Explicit expression for $r_{1,n}$

By 4.B.4, we have

$$J_2\left(3, \sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, T/\pi, \widetilde{K}_{3,n}, K_{4,n}, T, n\right)$$

$$\leq \frac{1.0253 \times 8\sqrt{2}}{2\pi} \left(\Gamma\left(3/2, \frac{\varepsilon}{4}\left(\frac{n}{K_{4,n}}\right)^{1/2}\right) - \Gamma\left(3/2, \frac{n}{2\widetilde{K}_{3,n}^2}\right)\right).$$

Moreover, we have

$$I_{3,3}(T) = \frac{1.0253|\lambda_{3,n}|}{6\pi\sqrt{n}} \bigg(\Gamma \Big(1, \varepsilon (n/K_{4,n})^{1/2} \Big) - \Gamma \Big(1, 2n/\widetilde{K}_{3,n}^2 \Big) \bigg).$$

Finally, we have

$$\frac{1.0253}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} R_{1,n}(u,\varepsilon) du = A_1(n) + \dots + A_7(n),$$

where

$$\begin{split} A_1(n) &:= \frac{1.0253}{2(1-3\varepsilon)^2 \pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} \frac{u^6}{24} \left(\frac{K_{4,n}}{n}\right)^{3/2} du, \\ A_2(n) &:= \frac{1.0253}{2(1-3\varepsilon)^2 \pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} \frac{u^8}{24^2} \left(\frac{K_{4,n}}{n}\right)^2 du, \\ A_3(n) &:= \frac{1.0253}{2(1-3\varepsilon)^2 \pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} \frac{u^5}{6} \left(\frac{K_{4,n}}{n}\right)^{5/4} du, \\ A_4(n) &:= \frac{1.0253}{2(1-3\varepsilon)^2 \pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} \frac{u^6}{36} \left(\frac{K_{4,n}}{n}\right)^{3/2} du, \\ A_5(n) &:= \frac{1.0253e_1(\varepsilon)}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} \frac{u^7}{72} \left(\frac{K_{4,n}}{n}\right)^{7/4} du, \\ A_6(n) &:= \frac{1.0253e_1(\varepsilon)}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} \frac{u^7 |\lambda_{3,n}| K_{4,n}}{6n^{3/2}} \left(\frac{1}{24} + \frac{P_{1,n}(\varepsilon)}{2(1-3\varepsilon)^2}\right)^2 du, \end{split}$$

where

$$P_{1,n}(\varepsilon) := \frac{144 + 48\varepsilon + 4\varepsilon^2 + \left\{96\sqrt{2\varepsilon} + 32\varepsilon + 16\sqrt{2}\varepsilon^{3/2}\right\} \mathbb{1}\left\{\exists i \in \{1, ..., n\} : \mathbb{E}[X_i^3] \neq 0\right\}}{576},$$
$$e_1(\varepsilon) := \exp\left(\varepsilon^2 \left(\frac{1}{6} + \frac{2P_{1,n}(\varepsilon)}{(1 - 3\varepsilon)^2}\right)\right).$$

Lemma 4.3. For any p > 0, $\int_0^{+\infty} u^p e^{-u^2/2} du = 2^{(p+1)/2} \Gamma((p+3)/2).$

Proof. We use the change of variable $v = u^2/2$, $u = \sqrt{2v}$, $dv = u du = \sqrt{2v}$, so that

$$\int_0^{+\infty} u^p e^{-u^2/2} du = \int_0^{+\infty} (2v)^{(p+1)/2} e^{-v} dv = 2^{(p+1)/2} \int_0^{+\infty} v^{(p+1)/2} e^{-v} dv,$$

and, by definition of $\Gamma(\cdot)$, this is equal to $2^{(p+1)/2}\Gamma((p+3)/2)$ as claimed.

Therefore, we can finally write that the $r_{1,n}$ of Equation (4.22) is upper bounded by

$$r_{1,n}^{\text{skew}} := \frac{a_9}{n^{5/4}} + \frac{a_{10}}{n^{6/4}} + \frac{a_{11}}{n^{7/4}} + \frac{a_1}{n^2} + \frac{a_2}{n^2} + \frac{a_3 \left(\Gamma(3/2, a_4 \sqrt{n}) - \Gamma(3/2, a_5 n) \right) + a_6 \left(\Gamma(1, a_7 \sqrt{n}) - \Gamma(1, a_8 n) \right)}{\sqrt{n}}, \qquad (4.23)$$

with

$$\begin{aligned} &-a_{1} := (14.1961 + 67.0415)\widetilde{K}_{3,n}^{4}/16\pi^{4} + (2^{12/2} \times \Gamma(12/2) \times 1.0253/\pi)K_{4,n}^{2} \times (1/(2 \times 24^{2})(1 - 3\varepsilon)^{2} + 1/24 + P_{1,n}(\varepsilon)/2(1 - 3\varepsilon)^{2}), \\ &-a_{2} := 4.3394|\lambda_{3,n}|\widetilde{K}_{3,n}^{3}/(8\pi^{3}), \\ &-a_{3} := 1.0253 \times 8\sqrt{2}K_{3,n}/6\pi, \\ &-a_{4} := \varepsilon/4\sqrt{K_{4,n}}, \\ &-a_{5} := 1/2\widetilde{K}_{3,n}^{2}, \\ &-a_{6} := 1.0253|\lambda_{3,n}|/6\pi, \\ &-a_{7} := \varepsilon/K_{4,n}^{1/2}, \\ &-a_{8} := 2/\widetilde{K}_{3,n}^{2}, \\ &-a_{9} := K_{4,n}^{5/4} \times 2^{9/2} \times \Gamma(9/2) \times 1.0253/12(1 - 3\varepsilon)^{2}\pi, \\ &-a_{10} := (1.0253/\pi) \times \left(2^{10/2} \times \Gamma(10/2) \times K_{4,n}^{3/2} \times (1/24 + 1/36)/2(1 - 3\varepsilon)^{2} + 2^{11/2} \times \Gamma(11/2) \times |\lambda_{3,n}| K_{4,n}/6 \times (1/24 + P_{1,n}(\varepsilon)/2(1 - 3\varepsilon)^{2})\right), \end{aligned}$$

The leading term of $r_{1,n}^{\text{skew}}$ decreases at rate $n^{5/4}$, hence the bound $r_{1,n} = O(n^{-5/4})$ in this setting "no-continuity *i.n.i.d.*" with possible skewness provided $K_{4,n} = O(1)$.

4.A.4 Proof of Theorem 4.1 under Assumption 4.2

We present and prove a more general result, Theorem 4.4, and choose $\varepsilon = 0.1$ to recover Theorem 4.1 under Assumption 4.2 (numerical computations of $e_3(0.1)$ give the upper bound $e_3(0.1) \leq 1.012$).

Theorem 4.4 (One-term Edgeworth expansion under Assumption 4.2). Let Assumption 4.2 hold. For every $\varepsilon \in (0, 1/3)$ and every $n \ge 2$

$$\Delta_{n,E} \leq \frac{0.1995\widetilde{K}_{3,n}}{\sqrt{n}} + \frac{1}{n} \left\{ 0.031\widetilde{K}_{3,n}^2 + 0.327K_{4,n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) + 0.054|\lambda_{3,n}|\widetilde{K}_{3,n} + 0.037e_3(\varepsilon)|\lambda_{3,n}|^2 \right\} + r_{1,n}, \quad (4.24)$$

where $r_{1,n}$ is given in Equation (4.30) and $e_3(\varepsilon) = e^{\varepsilon^2/6 + \varepsilon^2/(2(1-3\varepsilon))^2}$. We remark that $r_{1,n} = O(n^{-5/4})$ whenever $K_{4,n} = O(1)$. If $\mathbb{E}[X_n^3] = 0$ the upper bound reduces to

$$\frac{0.1995\widetilde{K}_{3,n}}{\sqrt{n}} + \frac{1}{n} \left\{ 0.031\widetilde{K}_{3,n}^2 + 0.327K_{4,n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) \right\} + r_{1,n},$$
(4.25)

with $r_{1,n} = O(n^{-2})$ when $K_{4,n} = O(1)$.

The overall method of proof is close to that in Section 4.A.3. In particular, the start of the proof is unchanged: we apply Lemma 4.1, choosing $t_0 = 1/\pi$, and Ω_1 is still controlled by Lemma 4.4. We keep with the choice $T = 2\pi\sqrt{n}/\tilde{K}_{3,n}$. The steps leading to an upper bound on Ω_2 in Section 4.A.3 remain valid as well so that we can write

$$\Omega_2(1/\pi, T) \le \frac{67.0415}{T^4} + \frac{1.2187}{T^2},\tag{4.26}$$

as in Equation (4.15). Thus, as described in Section 4.A.2, the treatment of Ω_1 and Ω_2 are common for Theorems 4.3 and 4.4. In contrast, as detailed below, the *i.i.d.* setting of Assumption 4.2 enables to improve the control of the third term Ω_3 coming from the smoothing inequality.

Improved bound on Ω_3

The control of Ω_3 can be refined under Assumption 4.2. We have

$$\Omega_3 \left(1/\pi, T, \lambda_{3,n} / \sqrt{n} \right) = 2 \int_0^{1/\pi} |\Psi(t)| \left| f_{S_n}(Tt) - e^{-(Tt)^2/2} \left(1 - \frac{\lambda_{3,n} i (Tt)^3}{6\sqrt{n}} \right) \right| dt$$

$$\leq I_{4,1}(T) + I_{4,2}(T) + I_{4,3}(T), \qquad (4.27)$$

where

$$\begin{split} I_{4,1}(T) &:= \frac{2}{T} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \left(1 - \frac{iu^3 \lambda_{3,n}}{6\sqrt{n}} \right) \right| du, \\ I_{4,2}(T) &:= \frac{2}{T} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T/\pi} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \right| du, \\ I_{4,3}(T) &:= \frac{2}{T} \frac{|\lambda_{3,n}|}{6\sqrt{n}} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T/\pi} |\Psi(u/T)| e^{-u^2/2} |u|^3 du. \end{split}$$

The integrand of $I_{4,1}(T)$ can be upper bounded thanks to Lemma 4.6. We obtain

$$I_{4,1}(T) \leq \frac{K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) J_1\left(4, 0, \sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, T\right) + \frac{e_{2,n}(\varepsilon)|\lambda_{3,n}|^2}{36n} J_1\left(6, 0, \sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, T\right) + \frac{2}{T} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} |\Psi(u/T)| e^{-u^2/2} R_{2,n}(u, \varepsilon) du \leq \frac{0.327 K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) + \frac{1.306 e_{2,n}(\varepsilon)|\lambda_{3,n}|^2}{36n} + \frac{1.0253}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} R_{2,n}(u, \varepsilon) du, \qquad (4.28)$$

using the expression of J_1 given in Equation (4.53) and the first inequality in (4.9). Moreover, we remark that

$$I_{4,3}(T) = \frac{|\lambda_{3,n}|}{3\sqrt{n}} J_1(3,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T/\pi,T)$$

As in Section 4.B.4, we can prove that $I_{4,3}(T)$ decays exponentially with n.

We finally control the term $I_{4,2}(T)$. Under the *i.i.d.* assumption, we can prove that, for every real t,

$$\left| f_{S_n}(t) - e^{-t^2/2} \right| \le \frac{K_{3,n}}{6\sqrt{n}} |t|^3 \exp\left(-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2}{2n}\right),$$

following the method of Section 4.A.3. Multiplying by $|\Psi(t)|$ and integrating this inequality, we get

$$I_{4,2}(T) \le \frac{K_{3,n}}{3\sqrt{n}} J_3(3, \sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, T/\pi, \widetilde{K}_{3,n}, T, n),$$
(4.29)

where

$$J_{3}(p,v,w,\widetilde{K}_{3,n},T,n) := \frac{1}{T} \int_{v}^{w} |\Psi(u/T)| u^{p} \exp\left(-\frac{u^{2}}{2}\left(1 - \frac{2\chi_{1}|u|\widetilde{K}_{3,n}}{\sqrt{n}} - \frac{1}{n}\right)\right) du.$$

Recalling that $T = 2\pi \sqrt{n}/\widetilde{K}_{3,n}$ we obtain

$$J_{3}(3,\sqrt{2\varepsilon(n/K_{4,n})^{1/4}},T/\pi,K_{3,n},T,n) = \frac{1}{T} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{2\sqrt{n}/\widetilde{K}_{3,n}} |\Psi(u/T)| u^{3} \exp\left(-\frac{u^{2}}{2}\left(1-\frac{2\chi_{1}|u|\widetilde{K}_{3,n}}{\sqrt{n}}-\frac{1}{n}\right)\right) du$$
$$\leq \frac{1}{T} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{2\sqrt{n}/\widetilde{K}_{3,n}} |\Psi(u/T)| u^{3} \exp\left(-\frac{u^{2}}{2}\left(1-4\chi_{1}-\frac{1}{n}\right)\right) du.$$

Note that $1 - 4\chi_1 - 1/n > 0.1$ as soon as $n \ge 2$. As in Section 4.B.4, we can prove that the latter term decays exponentially with n. The term $I_{4,2}(T)$ is thus negligible. Using the bound on $|\Psi(\cdot)|$ of Equation (4.9), we finally obtain

$$I_{4,2}(T) \le \frac{1.0253K_{3,n}}{6\pi\sqrt{n}} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{2\sqrt{n}} u^2 \exp\left(-\frac{t^2}{2} + \frac{\chi_1|t|^3\widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2}{2n}\right) du.$$

If $\mathbb{E}[X_n^3] = 0$, the bound (4.29) can be further improved to

$$I_{4,2}(T) \le \frac{K_{4,n}}{3n} J_3(4, \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}, T/\pi, \widetilde{K}_{3,n}, T, n)$$

This can be recovered using the same techniques as in Section 4.A.3, ensuring again that $I_{4,2}(T)$ decays exponentially fast to zero with n. Thanks to the control on $|\Psi(u/T)|$, if there is no skewness, we can upper bound $I_{4,2}(T)$ by

$$I_{4,2}(T) \le \frac{1.0253K_{4,n}}{6\pi n} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{2\sqrt{n}/\widetilde{K}_{3,n}} u^3 \exp\left(-\frac{t^2}{2} + \frac{\chi_1|t|^3\widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2}{2n}\right) du$$

Conclusion: end of the proof of Theorem 4.4

To conclude, we first use Equation (4.8), and manage all the terms separately. Ω_1 is bounded in Lemma 4.4, Ω_2 is bounded in Equation (4.26) and Ω_3 is bounded in Equation (4.28). We can claim that

$$\begin{split} \sup_{x \in \mathbb{R}} |\mathbb{P} \left(S_n \leq x \right) - G_n(x)| &\leq \Omega_1(t_0, T, |\lambda_{3,n}| / \sqrt{n}) + \Omega_2(t_0, T) + \Omega_3(t_0, T, \lambda_{3,n} / \sqrt{n}) \\ &\leq \frac{1.2533}{T} + \frac{0.3334 |\lambda_{3,n}|}{T\sqrt{n}} + \frac{14.1961}{T^4} + \frac{4.3394 |\lambda_{3,n}|}{T^3\sqrt{n}} \\ &\quad + \frac{67.0415}{T^4} + \frac{1.2187}{T^2} + I_{4,1}(T) + I_{4,2}(T) + I_{4,3}(T) \\ &\leq \frac{1.2533 \widetilde{K}_{3,n}}{2\pi\sqrt{n}} + \frac{0.3334 |\lambda_{3,n}| \widetilde{K}_{3,n}}{2\pi n} + \frac{1.2187 \widetilde{K}_{3,n}^2}{4\pi^2 n} \\ &\quad + \frac{0.327 K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) + \frac{1.306 e_3(\varepsilon) |\lambda_{3,n}|^2}{36n} + r_{1,n}, \end{split}$$

where

$$r_{1,n} := \frac{(14.1961 + 67.0415)\widetilde{K}_{3,n}^4}{2^4 \pi^4 n^2} + \frac{4.3394 |\lambda_{3,n}| \widetilde{K}_{3,n}^3}{8\pi^3 n^2} + \frac{K_{3,n}}{3\sqrt{n}} J_3(3,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, T/\pi, \widetilde{K}_{3,n}, T, n) + I_{4,3}(T) + \frac{1.306(e_{2,n}(\varepsilon) - e_3(\varepsilon)) |\lambda_{3,n}|^2}{36n} + \frac{1.0253}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} R_{2,n}(u,\varepsilon) du \quad (4.30)$$

The last term of Equation (4.30) is of order $n^{-3/2}$ given the definition of $R_{2,n}(u,\varepsilon)$. The quantity $(e_{2,n}(\varepsilon) - e_3(\varepsilon))/n$ is of the order $n^{-5/4}$ and is therefore dominant. Grouping terms together yields Equation (4.24).

When $\mathbb{E}[X_n^3] = 0$, we have $\lambda_{3,n} = 0$ which removes the corresponding terms. The first term on the second line in (4.30) becomes $\frac{K_{4,n}}{3n}J_3(4,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T/\pi,\widetilde{K}_{3,n},T,n)$. Furthermore, the leading terms in $r_{1,n}$ are $O(n^{-2})$ when $\lambda_{3,n} = 0$ given the definition of $R_{2,n}(u,\varepsilon)$.

4.A.5 Proof of Theorem 4.2 under Assumption 4.1

We use Theorem 4.5, proved below, with the choice $\varepsilon = 0.1$. We let $t_1^* := \theta_1^*/(2\pi)$ where θ_1^* is the unique root in $(0, 2\pi)$ of the equation $\theta^2 + 2\theta \sin(\theta) + 6(\cos(\theta) - 1) = 0$. Numerical approximations allow us to write $t_1^* \approx 0.64$.

Theorem 4.5 (Alternative one-term Edgeworth expansion under Assumption 4.1). Let Assumption 4.1 hold. For every $\varepsilon \in (0, 1/3)$ and every $n \ge 2$

$$\Delta_{n,E} \leq \frac{1}{n} \left\{ 0.327 K_{4,n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) + 0.037 e_{1,n}(\varepsilon) \lambda_{3,n}^2 \right\} \\ + \frac{1.0253}{\pi} \int_{a_n}^{b_n} \frac{|f_{S_n}(t)|}{t} dt + r_{2,n},$$
(4.31)

where $a_n = 2t_1^*\pi/\widetilde{K}_{3,n}$, $b_n = 4\pi^2 n/(t_1^*\widetilde{K}_{3,n}^2)$ and $r_{2,n}$ is given in Equation (4.35). When $K_{4,n} = O(1)$, we have $r_{2,n} = O(n^{-5/4})$. If $\mathbb{E}[X_i^3] = 0$ for every $i = 1, \ldots, n$, the upper bound becomes

$$\frac{0.327K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2}\right) + \frac{1.0253}{\pi} \int_{a_n}^{b_n} \frac{|f_{S_n}(t)|}{t} dt + r_{2,n},$$
(4.32)

where $r_{2,n} = O(n^{-3/2})$ when $K_{4,n} = O(1)$.

As before, we start the proof of Theorem 4.5 by applying Lemma 4.1 (with $t_0 = 1/\pi$) and use Lemma 4.4 to control Ω_1 as in the "no-continuity" setting. As explained in Section 4.A.2, in the "continuity" settings, the bounds on Ω_2 and Ω_3 are obtained in a different way, which is explained below. We choose $T = 16\pi^4 n^2 / \widetilde{K}_{3,n}^4$.

Bound on Ω_3

We decompose this term in five parts

$$\Omega_3(t_0, T, \lambda_{3,n}/\sqrt{n}) \le I_{5,1}(T) + I_{5,2}(T) + I_{5,3}(T) + I_{5,4}(T) + I_{5,5}(T),$$

where

$$\begin{split} I_{5,1}(T) &:= \frac{2}{T} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \left(1 - \frac{\lambda_{3,n} i u^3}{6\sqrt{n}} \right) \right| du, \\ I_{5,2}(T) &:= E_{1,n} \frac{|\lambda_{3,n}|}{3T\sqrt{n}} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T^{1/4}/\pi} |\Psi(u/T)| u^3 e^{-u^2/2} du, \\ I_{5,3}(T) &:= E_{1,n} \frac{2}{T} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T^{1/4}/\pi} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \right| du, \\ I_{5,4}(T) &:= E_{2,n} \frac{|\lambda_{3,n}|}{3T\sqrt{n}} \int_{T^{1/4}/\pi}^{T/\pi} |\Psi(u/T)| |u|^3 e^{-u^2/2} du, \\ I_{5,5}(T) &:= E_{2,n} \frac{2}{T} \int_{T^{1/4}/\pi}^{T/\pi} |\Psi(u/T)| \left| f_{S_n}(u) - e^{-u^2/2} \right| du, \end{split}$$

where $E_{1,n} := \mathbb{1}_{\{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4} < T^{1/4}/\pi\}}$ and $E_{2,n} := \mathbb{1}_{\{T^{1/4} < T\}}$. Note that if $T^{1/4} > T$ or $\sqrt{2\varepsilon}(n/K_{4,n})^{1/4} > T^{1/4}/\pi$, our bounds are still valid and can even be improved in the sense that the corresponding integrals can be removed.

Remarking that $I_{5,1}(T) = I_{3,1}(T)$, we can bound this term using Equation (4.17). We now turn to $I_{5,2}(T)$ and $I_{5,3}(T)$. Assume that $\sqrt{2\varepsilon}(n/K_{4,n})^{1/4} < T^{1/4}/\pi$, as the bound is trivially proved in the other case. We remark that $I_{5,2}(T)$ (resp. $I_{5,3}(T)$) can be bounded exactly as $I_{3,3}(T)$ in Section 4.A.3 (resp. as $I_{3,2}(T)$). Consequently,

$$I_{5,2}(T) \leq \frac{|\lambda_{3,n}|}{3\sqrt{n}} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T^{1/4}/\pi} \frac{1.0253}{2\pi} u^2 e^{-u^2/2} du$$

= $\frac{1.0253 |\lambda_{3,n}|}{3\pi\sqrt{2}\sqrt{n}} \Big(\Gamma\Big(3/2, \varepsilon(n/K_{4,n})^{1/2}\Big) - \Gamma\Big(3/2, T^{1/2}/2\pi^2\Big) \Big),$

and

$$\begin{split} I_{5,3}(T) &\leq \frac{2}{T} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T^{1/4}/\pi} |\Psi(u/T)| \frac{K_{3,n}}{6\sqrt{n}} |t|^3 \exp\left(-\frac{t^2}{2} + \frac{\chi_1 |t|^3 \widetilde{K}_{3,n}}{\sqrt{n}} + \frac{t^2 \sqrt{K_{4,n}}}{2\sqrt{n}}\right) du \\ &= \frac{K_{3,n}}{3\sqrt{n}} J_2 \Big(3, \sqrt{2\varepsilon}(n/K_{4,n})^{1/4}, T^{1/4}/\pi, \widetilde{K}_{3,n}, K_{4,n}, T, n\Big), \end{split}$$

and these terms decrease to zero exponentially fast when n goes to infinity.

Note that whenever $\mathbb{E}[X_i^3] = 0$ for all i = 1, ..., n, the improvements detailed in Section 4.A.3 can be used as well, resulting in a tighter bound on $I_{5,2}$ where the factor $K_{3,n}/\sqrt{n}$ is replaced by $K_{4,n}/n$ and the first argument of J_2 becomes 4 instead of 3. We have to deal with $I_{5,4}(T)$ and $I_{5,5}(T)$. We assume that $T^{1/4} < T$, otherwise there is nothing to prove. Obviously $I_{5,4}(T)$ can be bounded in a similar vein as $I_{5,2}(T)$

$$I_{5,4}(T) = \frac{1.0253 |\lambda_{3,n}|}{3\pi\sqrt{2}\sqrt{n}} \Big(\Gamma \Big(3/2 \,, \, T^{1/2}/2\pi^2 \Big) - \Gamma \Big(3/2 \,, \, T^2/2\pi^2 \Big) \Big),$$

and it converges exponentially fast to zero.

To control $I_{5,5}(T)$, we write

$$I_{5,5}(T) \le J_3(T) + J_4(T) + J_5(T)$$

where

$$J_{3}(T) := \frac{2}{T} \int_{T^{1/4}/\pi}^{t_{1}^{*}T^{1/4}} |\Psi(u/T)| |f_{S_{n}}(u)| \, du = \frac{2}{T^{3/4}} \int_{1/\pi}^{t_{1}^{*}} |\Psi(v/T^{3/4})| \left| f_{S_{n}}(T^{1/4}v) \right| \, dv,$$

$$J_{4}(T) := \mathbb{1}_{\left\{ t_{1}^{*}T^{1/4} < T/\pi \right\}} \frac{2}{T} \int_{t_{1}^{*}T^{1/4}}^{T/\pi} |\Psi(u/T)| |f_{S_{n}}(u)| \, du,$$

$$J_{5}(T) := \frac{2}{T} \int_{T^{1/4}/\pi}^{T/\pi} |\Psi(u/T)| e^{-u^{2}/2} \, du.$$

By Lemma 4.2 and our choice of T, we know $|f_{S_n}(T^{1/4}v)|$ can be upper bounded by $e^{-\frac{T^{1/2}v^2}{2}(1-4\pi\chi_1|v|)}$ when $v \in [1/\pi, t_1^*]$. We get (using $1 - 4\pi\chi_1 t_1^* > 0$)

$$\begin{aligned} J_{3}(T) &\leq \frac{2}{T^{3/4}} \int_{1/\pi}^{t_{1}^{*}} |\Psi(v/T^{3/4})| e^{-\frac{T^{1/2}v^{2}}{2}(1-4\pi\chi_{1}|v|)} dv \\ &\leq \frac{1.0253}{\pi} \int_{1/\pi}^{t_{1}^{*}} v^{-1} e^{-\frac{T^{1/2}v^{2}}{2}\left(1-4\pi\chi_{1}t_{1}^{*}\right)} dv \\ &= \frac{1.0253}{2\pi} \int_{T^{1/2}(1-4\pi\chi_{1}t_{1}^{*})/(2\pi^{2})}^{t_{1}^{*2}T^{1/2}(1-4\pi\chi_{1}t_{1}^{*})/2} u^{-1} e^{-u} du \\ &= \frac{1.0253}{2\pi} \left(\Gamma \left(T^{1/2}(1-4\pi\chi_{1}t_{1}^{*})/(2\pi^{2}), 0 \right) - \Gamma \left(t_{1}^{*2}T^{1/2}(1-4\pi\chi_{1}t_{1}^{*})/2, 0 \right) \right). \end{aligned}$$

As a result, we conclude that $J_3(T)$ decreases to zero exponentially fast with n.

To control $J_4(T)$, we use the properties of $u \mapsto \Psi(u)$ to write

$$J_4(T) \le \mathbb{1}_{\left\{t_1^* T^{1/4} < T/\pi\right\}} \frac{1.0253}{\pi} \int_{t_1^* T^{1/4}}^{T/\pi} u^{-1} |f_{S_n}(u)| du$$

To upper bound $J_5(T)$, we can reason as for $J_3(T)$ to conclude that this term converges to zero exponentially fast.

As a result, we conclude

$$\Omega_{3}(1/\pi, T, \lambda_{3,n}/\sqrt{n}) \leq I_{5,1}(T) + I_{5,2}(T) + I_{5,3}(T) + I_{5,4}(T) + I_{5,5}(T) \\
\leq \frac{0.327K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^{2}}\right) + \frac{1.306e_{1,n}(\varepsilon)}{36} \frac{|\lambda_{3,n}|^{2}}{n} \\
+ \mathbb{1}_{\left\{t_{1}^{*}T^{1/4} < T/\pi\right\}} \frac{1.0253}{\pi} \int_{t_{1}^{*}T^{1/4}}^{T/\pi} u^{-1} |f_{S_{n}}(u)| du + O(n^{-5/4}). \quad (4.33)$$

Conclusion: end of the proof of Theorem 4.5

We can write

$$\Omega_2(1/\pi, T) \le \frac{1.0253}{\pi} \int_{T/\pi}^T u^{-1} |f_{S_n}(u)| du.$$
(4.34)

To sum up, we first use Equation (4.8), and manage all the terms separately. Ω_1 is bounded in Lemma 4.4, Ω_2 is bounded in Equation (4.34) and Ω_3 is bounded in Equation (4.33). Given the definitions of a_n and b_n , we conclude

$$\begin{split} \sup_{x \in \mathbb{R}} |\mathbb{P} \left(S_n \leq x \right) - G_n(x) | &\leq \Omega_1(t_0, T, |\lambda_{3,n}| / \sqrt{n}) + \Omega_2(t_0, T) + \Omega_3(t_0, T, \lambda_{3,n} / \sqrt{n}) \\ &\leq \frac{1.2533}{T} + \frac{0.3334 |\lambda_{3,n}|}{T\sqrt{n}} + \frac{14.1961}{T^4} + \frac{4.3394 |\lambda_{3,n}|}{T^3\sqrt{n}} \\ &\quad + \frac{1.0253}{\pi} \int_{T/\pi}^T u^{-1} |f_{S_n}(u)| du + I_{5,1}(T) + I_{5,2}(T) + I_{5,3}(T) + I_{5,4}(T) + I_{5,5}(T) \\ &\leq \frac{0.327 K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1 - 3\varepsilon)^2} \right) + \frac{1.306 e_{1,n}(\varepsilon) |\lambda_{3,n}|^2}{36n} \\ &\quad + \frac{1.0253}{\pi} \int_{a_n}^{b_n} u^{-1} |f_{S_n}(u)| du + r_{2,n}, \end{split}$$

where

$$r_{2,n} := \frac{1.2533\widetilde{K}_{3,n}^4}{16\pi^4 n^2} + \frac{0.3334|\lambda_{3,n}|\widetilde{K}_{3,n}^4}{16\pi^4 n^{5/2}} + \frac{14.1961\widetilde{K}_{3,n}^{16}}{16^4 \pi^{16} n^8} + \frac{4.3394|\lambda_{3,n}|\widetilde{K}_{3,n}^{12}}{16^3 \pi^{12} n^6 \sqrt{n}} \\ + I_{5,2}(T) + I_{5,4}(T) + J_3(T) + J_5(T) + \frac{K_{3,n}}{3\sqrt{n}} J_2 \Big(3, \sqrt{2\varepsilon}/(nK_{4,n})^{1/4}, T^{1/4}/\pi, \widetilde{K}_{3,n}, K_{4,n}, T, n\Big) \\ + \frac{1.0253}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} R_{1,n}(u, \varepsilon) du.$$

$$(4.35)$$

All terms but the last one in the definition of $r_{2,n}$ are at most of order n^{-2} . As explained in Section 4.A.3, the last term is of order $O(n^{-5/4})$ so that $r_{2,n} = O(n^{-5/4})$.

When $\mathbb{E}[X_i^3] = 0$ for every i = 1, ..., n, we have $\lambda_{3,n} = 0$, which removes the corresponding terms. The final term on the second line in (4.35) becomes $\frac{K_{4,n}}{3n}J_2\left(4,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T^{1/4}/\pi,\widetilde{K}_{3,n},K_{4,n},T,n\right)$. Under the new assumption, the dominant term in $r_{2,n}$ has the rate $n^{-3/2}$ which implies $r_{2,n} = O(n^{-3/2})$.

4.A.6 Proof of Theorem 4.2 under Assumption 4.2

We use Theorem 4.6, proved below, with the choice $\varepsilon = 0.1$.

Theorem 4.6 (Alternative one-term Edgeworth expansion under Assumption 4.2). Let Assumption 4.2 hold. For every $\varepsilon \in (0, 1/3)$ and every $n \ge 1$

$$\Delta_{n,E} \leq \frac{1}{n} \left\{ 0.327 K_{4,n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2} \right) + 0.037 e_3(\varepsilon) \lambda_{3,n}^2 \right\} \\ + \frac{1.0253}{\pi} \int_{a_n}^{b_n} \frac{|f_{S_n}(t)|}{t} dt + r_{2,n},$$
(4.36)

where $a_n = 2t_1^* \pi / \widetilde{K}_{3,n}$, $b_n = 4\pi^2 n / (t_1^* \widetilde{K}_{3,n}^2)$ and $r_{2,n}$ is given in Equation (4.40). We have $r_{2,n} = O(n^{-5/4})$ as soon as $K_{4,n} = O(1)$. If $\mathbb{E}[X_n^3] = 0$, the upper bound becomes

$$\frac{0.327K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^2}\right) + \frac{1.0253}{\pi} \int_{a_n}^{b_n} \frac{|f_{S_n}(t)|}{t} dt + r_{2,n},$$
(4.37)

where $r_{2,n} = O(n^{-2})$ when $K_{4,n} = O(1)$.

The proof of Theorem 4.6 is very similar to that of Theorem 4.5. We start by applying Lemma 4.1 (with $t_0 = 1/\pi$) and use Lemma 4.4 to control Ω_1 . There remains to bound Ω_2 and Ω_3 . We choose $T = 16\pi^4 n^2 / \widetilde{K}_{3,n}^4$ again.

Bound on Ω_3

As in the proof of Theorem 4.5, we write

$$\Omega_3(t_0, T, \lambda_{3,n}/\sqrt{n}) \le I_{5,1}(T) + I_{5,2}(T) + I_{5,3}(T) + I_{5,4}(T) + I_{5,5}(T),$$

where the terms in the upper bound are defined in Section 4.A.5.

The terms $I_{5,2}(T)$ and $I_{5,4}(T)$ are controlled similarly as in the *i.n.i.d.* case of Section 4.A.5. We have $I_{5,1}(T) = I_{4,1}(T)$ so that we can use the upper bound in (4.28). The term $I_{5,3}(T)$ can be controlled as $I_{4,2}(T)$. We upper bound $I_{5,5}(T)$ as in Section 4.A.5

$$I_{5,5}(T) \le J_3(T) + J_4(T) + J_5(T).$$

The proof that $J_3(T)$ and $J_5(T)$ decrease exponentially fast to zero with n is still valid. We finally obtain

We finally obtain

$$\Omega_{3}(1/\pi, T, \lambda_{3,n}/\sqrt{n}) \leq I_{5,1}(T) + I_{5,2}(T) + I_{5,3}(T) + I_{5,4}(T) + I_{5,5}(T) \\
\leq \frac{0.327K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1-3\varepsilon)^{2}}\right) + \frac{1.306e_{3}(\varepsilon)}{36} \frac{|\lambda_{3,n}|^{2}}{n} \\
+ \mathbb{1}_{\left\{t_{1}^{*}T^{1/4} < T/\pi\right\}} \frac{1.0253}{\pi} \int_{t_{1}^{*}T^{1/4}}^{T/\pi} u^{-1} |f_{S_{n}}(u)| du.$$
(4.38)

Conclusion: end of the proof of Theorem 4.6

We have

$$\Omega_2(1/\pi, T) \le \frac{1.0253}{\pi} \int_{T/\pi}^T u^{-1} |f_{S_n}(u)| du.$$
(4.39)

To sum up, we first use Equation (4.8), and manage all the terms separately. Ω_1 is bounded in Lemma 4.4, Ω_2 is bounded in Equation (4.39) and Ω_3 is bounded in

Equation (4.38). Using the definitions of a_n and b_n , we conclude

$$\begin{split} \sup_{x \in \mathbb{R}} |\mathbb{P} \left(S_n \leq x \right) - G_n(x) | &\leq \Omega_1(t_0, T, |\lambda_{3,n}| / \sqrt{n}) + \Omega_2(t_0, T) + \Omega_3(t_0, T, \lambda_{3,n} / \sqrt{n}) \\ &\leq \frac{1.2533}{T} + \frac{0.3334 |\lambda_{3,n}|}{T\sqrt{n}} + \frac{14.1961}{T^4} + \frac{4.3394 |\lambda_{3,n}|}{T^3\sqrt{n}} \\ &\quad + \frac{1.0253}{\pi} \int_{T/\pi}^T u^{-1} |f_{S_n}(u)| du \\ &\quad + I_{5,1}(T) + I_{5,2}(T) + I_{5,3}(T) + I_{5,4}(T) + I_{5,5}(T) \\ &\leq \frac{0.327 K_{4,n}}{n} \left(\frac{1}{12} + \frac{1}{4(1 - 3\varepsilon)^2} \right) + \frac{1.306 e_3(\varepsilon) |\lambda_{3,n}|^2}{36n} \\ &\quad + \frac{1.0253}{\pi} \int_{a_n}^{b_n} u^{-1} |f_{S_n}(u)| du + r_{2,n}, \end{split}$$

where

$$r_{2,n} := \frac{1.2533\widetilde{K}_{3,n}^4}{16\pi^4 n^2} + \frac{0.3334|\lambda_{3,n}|\widetilde{K}_{3,n}^4}{16\pi^4 n^{5/2}} + \frac{14.1961\widetilde{K}_{3,n}^{16}}{16^4 \pi^{16} n^8} + \frac{4.3394|\lambda_{3,n}|\widetilde{K}_{3,n}^{12}}{16^3 \pi^{12} n^6 \sqrt{n}} + I_{5,2}(T) + I_{5,4}(T) + J_3(T) + J_5(T) + \frac{K_{3,n}}{3\sqrt{n}} J_3(3,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T^{1/4}/\pi,\widetilde{K}_{3,n},T,n) + \frac{1.306(e_{2,n}(\varepsilon) - e_3(\varepsilon))|\lambda_{3,n}|^2}{36n} + \frac{1.0253}{\pi} \int_0^{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}} u e^{-u^2/2} R_{2,n}(u,\varepsilon) du.$$
(4.40)

All terms but the last one in the definition of $r_{2,n}$ are at most of order n^{-2} . As explained in Section 4.A.4, the last term is of order $O(n^{-5/4})$ so that $r_{2,n} = O(n^{-5/4})$.

When $\mathbb{E}[X_n^3] = 0$, we have $\lambda_{3,n} = 0$ which removes the corresponding terms. The final term on the second line in (4.40) becomes $\frac{K_{4,n}}{3n}J_3(4,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T^{1/4}/\pi,\widetilde{K}_{3,n},T,n)$. Under the new assumption, the dominant term in $r_{2,n}$ has the rate n^{-2} which implies $r_{2,n} = O(n^{-2})$.

Appendix 4.B Technical lemmas

4.B.1 Control of the term Ω_1

Lemma 4.4. For every T > 0, we have

$$\Omega_1(1/\pi, T, |\lambda_{3,n}|/\sqrt{n}) \le \frac{1.2533}{T} + \frac{0.3334|\lambda_{3,n}|}{T\sqrt{n}} + \frac{14.1961}{T^4} + \frac{4.3394|\lambda_{3,n}|}{T^3\sqrt{n}}.$$
 (4.41)

Proof. With the choice $t_0 = 1/\pi$, the function $\Omega_1(1/\pi, T, v)$ can be decomposed as

$$\Omega_1(t_0, T, v) := \frac{I_{1,1}(T)}{T} + v \times \frac{I_{1,2}(T)}{T} + \frac{I_{1,3}(T)}{T^4} + v \times \frac{I_{1,4}(T)}{T^3}$$

where

$$\begin{split} I_{1,1}(T) &:= T \int_0^{1/\pi} \left| 2\Psi(t) - \frac{i}{\pi t} \right| e^{-(Tt)^2/2} dt, \\ I_{1,2}(T) &:= T^4 \int_0^{1/\pi} \left| 2\Psi(t) - \frac{i}{\pi t} \right| e^{-(Tt)^2/2} \frac{t^3}{6} dt, \\ I_{1,3}(T) &:= T^4 \frac{1}{\pi} \int_{1/\pi}^{+\infty} \frac{e^{-(Tt)^2/2}}{t} dt = \frac{T^4}{2\pi} \Gamma\left(0, \frac{T^2}{2\pi^2}\right), \\ I_{1,4}(T) &:= T^6 \frac{1}{\pi} \int_{1/\pi}^{+\infty} e^{-(Tt)^2/2} \frac{t^2}{6} dt = \frac{T^3}{3\sqrt{2\pi}} \int_{T^2/(2\pi^2)}^{+\infty} e^{-u} \sqrt{u} du = \frac{T^3}{3\sqrt{2\pi}} \Gamma\left(\frac{3}{2}, \frac{T^2}{2\pi^2}\right). \end{split}$$

On the last two lines, we used the change of variable $u = (tT)^2/2$ and the incomplete Gamma function $\Gamma(a, x) := \int_x^{+\infty} u^{a-1} e^{-u} du$ which can be computed numberically using the package expirt (Goulet, 2016) in R. We estimate numerically the first two integrals using the R package cubature (Narasimhan et al., 2020) and optimize using the optimize function with the L-BFGS-B method, we find the following upper bounds:⁷

$$\sup_{T \ge 0} I_{1,1}(T) \le 1.2533, \qquad \qquad \sup_{T \ge 0} I_{1,2}(T) \le 0.3334, \\
\sup_{T \ge 0} I_{1,3}(T) \le 14.1961, \qquad \qquad \sup_{T \ge 0} I_{1,4}(T) \le 4.3394,$$

which finishes the proof.

Note that the first term on the right-hand side of (4.41) is of leading order as soon as $|\lambda_{3,n}|/\sqrt{n} = o(1)$ and T = T(n) = o(1). Our approach is related to the one used in Shevtsova (2012), except that we do not upper bound Ω_1 analytically, which allows us to get a sharper control on this term. To further highlight the gains from using numerical approximations instead of direct analytical upper bounds, we remark that from $|\Psi(t) - \frac{i}{2\pi t}| \leq \frac{1}{2} \left(1 - |t| + \frac{\pi^2 t^2}{18}\right)$ and some integration steps, we get

$$\begin{split} I_{1,1}(T) &\leq T \int_0^{1/\pi} \left(1 - |t| + \frac{\pi^2 t^2}{18} \right) e^{-(Tt)^2/2} dt \\ &= \sqrt{2\pi} \left(\Phi(T/\pi) - \frac{1}{2} \right) + \frac{1}{T} \left(e^{-(T/\pi)^2/2} - 1 \right) + \frac{\pi^{5/2}}{9\sqrt{2}T^2} \mathbb{E}_{U \sim \mathcal{N}(0,1)} [U^2 \mathbb{1} \left\{ 0 \leq U \leq T/\pi \right\}] \\ &\leq \sqrt{2\pi} + \frac{1}{T} \left(e^{-T^2/(2\pi^2)} - 1 \right) + \frac{\pi^{5/2}}{9\sqrt{2}T^2}, \end{split}$$

⁷We would be happy to share the code upon request.

whose main term is approximately twice as large as the numerical bound 1.2533 that we obtained before.

4.B.2 Control of the residual term in an Edgeworth expansion under Assumption 4.1

For $\varepsilon \in (0, 1/3)$ and $t \ge 0$, let us define the following quantities:

$$\begin{split} R_{1,n}(t,\varepsilon) &:= \frac{U_{1,1,n}(t) + U_{1,2,n}(t)}{2(1-3\varepsilon)^2} \\ &+ e_1(\varepsilon) \left(\frac{t^8 K_{4,n}^2}{2n^2} \left(\frac{1}{24} + \frac{P_{1,n}(\varepsilon)}{2(1-3\varepsilon)^2} \right)^2 + \frac{|t|^7 |\lambda_{3,n}| K_{4,n}}{6n^{3/2}} \left(\frac{1}{24} + \frac{P_{1,n}(\varepsilon)}{2(1-3\varepsilon)^2} \right) \right), \\ P_{1,n}(\varepsilon) &:= \frac{144 + 48\varepsilon + 4\varepsilon^2 + \left\{ 96\sqrt{2\varepsilon} + 32\varepsilon + 16\sqrt{2}\varepsilon^{3/2} \right\} \mathbf{1} \left\{ \exists i \in \{1, ..., n\} : \mathbb{E}[X_i^3] \neq 0 \right\}}{576}, \\ e_1(\varepsilon) &:= \exp\left(\varepsilon^2 \left(\frac{1}{6} + \frac{2P_{1,n}(\varepsilon)}{(1-3\varepsilon)^2} \right) \right), \\ U_{1,1,n}(t) &:= \frac{t^6}{24} \left(\frac{K_{4,n}}{n} \right)^{3/2} + \frac{t^8}{24^2} \left(\frac{K_{4,n}}{n} \right)^2, \\ U_{1,2,n}(t) &:= \left(\frac{|t|^5}{6} \left(\frac{K_{4,n}}{n} \right)^{5/4} + \frac{t^6}{36} \left(\frac{K_{4,n}}{n} \right)^{3/2} + \frac{|t|^7}{72} \left(\frac{K_{4,n}}{n} \right)^{7/4} \right) \mathbf{1} \left\{ \exists i \in \{1, ..., n\} : \mathbb{E}[X_i^3] \neq 0 \right\}. \end{aligned}$$

$$(4.43)$$

We want to show the following lemma:

Lemma 4.5. Under Assumption 4.1, for every $\varepsilon \in (0, 1/3)$ and t such that $|t| \leq \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}$, we have

$$\left| f_{S_n}(t) - e^{-\frac{t^2}{2}} \left(1 - \frac{it^3 \lambda_{3,n}}{6\sqrt{n}} \right) \right| \le e^{-t^2/2} \left\{ \frac{t^4 K_{4,n}}{8n} \left(\frac{1}{3} + \frac{1}{(1-3\varepsilon)^2} \right) + \frac{e_1(\varepsilon)|t|^6 |\lambda_{3,n}|^2}{72n} + R_{1,n}(t,\varepsilon) \right\}.$$

Proof of Lemma 4.5: Remember that $\gamma_j := \mathbb{E}[X_j^4], \sigma_j := \sqrt{\mathbb{E}[X_j^2]}, B_n := \sqrt{\sum_{i=1}^n \mathbb{E}[X_i^2]}$ and $K_{4,n} := n^{-1} \sum_{i=1}^n \mathbb{E}[X_i^4] / (n^{-1}B_n^2)^2$. Applying Cauchy-Schwartz inequality, we get

$$\max_{1 \le j \le n} \sigma_j^2 \le \max_{1 \le j \le n} \gamma_j^{1/2} \le \left(\sum_{j=1}^n \gamma_j\right)^{1/2} = B_n^2 (K_{4,n}/n)^{1/2}, \tag{4.44}$$

$$\max_{1 \le j \le n} \mathbb{E}[|X_j|^3] \le \max_{1 \le j \le n} \gamma_j^{3/4} \le \left(\sum_{j=1}^n \gamma_j\right)^{3/4} = B_n^3 (K_{4,n}/n)^{3/4}, \tag{4.45}$$

and

$$\max_{1 \le j \le n} \gamma_j \le \sum_{j=1}^n \gamma_j = B_n^4 K_{4,n}/n.$$
(4.46)

Combining (4.44), (4.45) and (4.46), we observe that for every $\varepsilon \in (0, 1)$ and t such that $|t| \leq \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}$,

$$\max_{1 \le j \le n} \left\{ \frac{\sigma_j^2 t^2}{2B_n^2} + \frac{\mathbb{E}[|X_j|^3] \times |t|^3}{6B_n^3} + \frac{\gamma_j t^4}{24B_n^4} \right\} \le 3\varepsilon.$$
(4.47)
As we assume that X_j has a moment of order four for every j = 1, ..., n, the characteristic functions $(f_{P_{X_j}})_{j=1,...,n}$ are four times differentiable on \mathbb{R} . Applying a Taylor-Lagrange expansion, we get the existence of a complex number $\theta_{1,j,n}(t)$ such that $|\theta_{1,j,n}(t)| \leq 1$ and

$$U_{j,n}(t) := f_{P_{X_j}}(t/B_n) - 1 = -\frac{\sigma_j^2 t^2}{2B_n^2} - \frac{i\mathbb{E}[X_j^3] t^3}{6B_n^3} + \frac{\theta_{1,j,n}(t)\gamma_j t^4}{24B_n^4},$$

for every $t \in \mathbb{R}$ and j = 1, ..., n. Let log stand for the principal branch of the complex logarithm function. For every $\varepsilon \in (0, 1/3)$ and t such that $|t| \leq \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}$, Equation (4.47) shows that $|U_{j,n}(t)| \leq 3\varepsilon < 1$, so that we can use another Taylor-Lagrange expansion. This ensures existence of a complex number $\theta_{2,j,n}(t)$ such that $|\theta_{2,j,n}(t)| \leq 1$ and

$$\log(f_{P_{X_j}}(t/B_n)) = \log(1 + U_{j,n}(t)) = U_{j,n}(t) - \frac{U_{j,n}(t)^2}{2(1 + \theta_{2,j,n}(t)U_{j,n}(t))^2}$$

Summing over j = 1, ..., n and exponentiating, we can claim that under the same conditions on t and ε ,

$$f_{S_n}(t) = \exp\left(-\frac{t^2}{2} - \frac{it^3\lambda_{3,n}}{6\sqrt{n}} + t^4\sum_{j=1}^n \frac{\theta_{1,j,n}(t)\gamma_j}{24B_n^4} - \sum_{j=1}^n \frac{U_{j,n}(t)^2}{2(1+\theta_{2,j,n}(t)U_{j,n}(t))^2}\right).$$

A third Taylor-Lagrange expansion guarantees existence of a complex number $\theta_{3,n}(t)$ with modulus at most $\exp\left(\frac{t^4K_{4,n}}{24n} + \sum_{j=1}^n \frac{|U_{j,n}(t)|^2}{2|1+\theta_{2,j,n}(t)U_{j,n}(t)|^2}\right)$ such that

$$f_{S_n}(t) = e^{-t^2/2} \left(1 - \frac{it^3\lambda_{3,n}}{6\sqrt{n}} + t^4 \sum_{j=1}^n \frac{\theta_{1,j,n}(t)\gamma_j}{24B_n^4} - \sum_{j=1}^n \frac{U_{j,n}(t)^2}{2(1+\theta_{2,j,n}(t)U_{j,n}(t))^2} + \frac{\theta_{3,n}(t)}{2} \left(-\frac{it^3\lambda_{3,n}}{6\sqrt{n}} + t^4 \sum_{j=1}^n \frac{\theta_{1,j,n}(t)\gamma_j}{24B_n^4} - \sum_{j=1}^n \frac{U_{j,n}(t)^2}{2(1+\theta_{2,j,n}(t)U_{j,n}(t))^2} \right)^2 \right).$$

Using the triangle inequality and its reverse version, as well as the restriction on $|t| \leq \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}$, we can write

$$\left| f_{S_n}(t) - e^{-t^2/2} \left(1 - \frac{it^3 \lambda_{3,n}}{6\sqrt{n}} \right) \right| \le e^{-t^2/2} \times \left(\frac{t^4 K_{4,n}}{24n} + \frac{1}{2(1-3\varepsilon)^2} \sum_{j=1}^n |U_{j,n}(t)|^2 + \frac{1}{2} \exp\left(\frac{\varepsilon^2}{6} + \frac{1}{2(1-3\varepsilon)^2} \sum_{j=1}^n |U_{j,n}(t)|^2 \right) \times \left(\frac{|t|^3 |\lambda_{3,n}|}{6\sqrt{n}} + \frac{t^4 K_{4,n}}{24n} + \frac{1}{2(1-3\varepsilon)^2} \sum_{j=1}^n |U_{j,n}(t)|^2 \right)^2 \right)$$

$$(4.48)$$

We now control $\sum_{j=1}^{n} |U_{j,n}(t)|^2$. We first expand the squares, giving the decomposition

$$\sum_{j=1}^{n} |U_{j,n}(t)|^{2} = \frac{t^{4} \sum_{j=1}^{n} \sigma_{j}^{4}}{4B_{n}^{4}} + \frac{t^{6} \sum_{j=1}^{n} |\mathbb{E}[X_{j}^{3}]|^{2}}{36B_{n}^{6}} + \frac{t^{8} \sum_{j=1}^{n} \gamma_{j}^{2}}{24^{2}B_{n}^{8}} + \frac{|t|^{5} \sum_{j=1}^{n} \sigma_{j}^{2}|\mathbb{E}[X_{j}^{3}]|}{6B_{n}^{5}} + \frac{t^{6} \sum_{j=1}^{n} \sigma_{j}^{2}\gamma_{j}}{24B_{n}^{6}} + \frac{|t|^{7} \sum_{j=1}^{n} |\mathbb{E}[X_{j}^{3}]|\gamma_{j}}{72B_{n}^{7}}.$$
 (4.49)

Using Equations (4.44)-(4.46), we can bound the right-hand side of Equation (4.49) in the following manner

$$\frac{t^4 \sum_{j=1}^n \sigma_j^4}{4B_n^4} \le \frac{t^4 K_{4,n}}{4n},$$

$$\frac{t^6 \sum_{j=1}^n \sigma_j^2 \gamma_j}{24B_n^6} + \frac{t^8 \sum_{j=1}^n \gamma_j^2}{24^2 B_n^8} \le \frac{t^6}{24} \left(\frac{K_{4,n}}{n}\right)^{3/2} + \frac{t^8}{24^2} \left(\frac{K_{4,n}}{n}\right)^2 =: U_{1,1,n}(t),$$

and

$$\frac{|t|^{5} \sum_{j=1}^{n} \sigma_{j}^{2} |\mathbb{E}[X_{j}^{3}]|}{6B_{n}^{5}} + \frac{t^{6} \sum_{j=1}^{n} |\mathbb{E}[X_{j}^{3}]|^{2}}{36B_{n}^{6}} + \frac{|t|^{7} \sum_{j=1}^{n} |\mathbb{E}[X_{j}^{3}]|\gamma_{j}}{72B_{n}^{7}} \\
\leq \left(\frac{|t|^{5}}{6} \left(\frac{K_{4,n}}{n}\right)^{5/4} + \frac{t^{6}}{36} \left(\frac{K_{4,n}}{n}\right)^{3/2} + \frac{|t|^{7}}{72} \left(\frac{K_{4,n}}{n}\right)^{7/4}\right) \mathbb{1}\left\{\exists i \in \{1,...,n\} : \mathbb{E}[X_{i}^{3}] \neq 0\right\} \\
=: U_{1,2,n}(t). \tag{4.50}$$

Moreover, we have $\sum_{j=1}^{n} U_{j,n}(t)^2 \leq \frac{t^4 K_{4,n}}{n} P_{1,n}(\varepsilon)$ under our conditions on ε and t. Combining Equation (4.48), the decomposition (4.49) and the previous three bounds, and grouping similar terms together, we conclude that for every $\varepsilon \in (0, 1/3)$ and t such that $|t| \leq \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}$,

$$\begin{aligned} \left| f_{S_n}(t) - e^{-\frac{t^2}{2}} \left(1 - \frac{it^3 \lambda_{3,n}}{6\sqrt{n}} \right) \right| \\ &\leq e^{-t^2/2} \Biggl\{ \frac{t^4 K_{4,n}}{8n} \left(\frac{1}{3} + \frac{1}{(1-3\varepsilon)^2} \right) + \frac{e_{1,n}(\varepsilon)|t|^6 |\lambda_{3,n}|^2}{72n} + \frac{U_{1,1,n}(t) + U_{1,2,n}(t)}{2(1-3\varepsilon)^2} \\ &+ e_1(\varepsilon) \left(\frac{t^8 K_{4,n}^2}{2n^2} \left(\frac{1}{24} + \frac{P_{1,n}(\varepsilon)}{2(1-3\varepsilon)^2} \right)^2 + \frac{|t|^7 |\lambda_{3,n}| K_{4,n}}{6n^{3/2}} \left(\frac{1}{24} + \frac{P_{1,n}(\varepsilon)}{2(1-3\varepsilon)^2} \right) \Biggr) \Biggr\}, \end{aligned}$$

where $e_1(\varepsilon) := \exp\left(\varepsilon^2\left(\frac{1}{6} + \frac{2P_{1,n}(\varepsilon)}{(1-3\varepsilon)^2}\right)\right)$. Combining this with the definition of $R_{1,n}(t,\varepsilon)$ finishes the proof. \Box

4.B.3 Control of the residual term in an Edgeworth expansion under Assumption 4.2

Lemma 4.5 can be improved in the *i.i.d.* framework. To do so, we introduce analogues of $R_{1,n}(t,\varepsilon)$, $P_{1,n}(\varepsilon)$, $e_{2,n}(\varepsilon)$ and $U_{1,2,n}(t)$ defined by

$$\begin{split} R_{2,n}(t,\varepsilon) &:= \frac{U_{2,2,n}(t)}{2(1-3\varepsilon)^2} + e_{2,n}(\varepsilon) \left(\frac{t^8}{8n^2} \left(\frac{K_{4,n}}{12} + \frac{P_{2,n}(\varepsilon)}{(1-3\varepsilon)^2} \right)^2 \right. \\ &+ \frac{|t|^7 |\lambda_{3,n}|}{12n^{3/2}} \left(\frac{K_{4,n}}{12} + \frac{P_{2,n}(\varepsilon)}{(1-3\varepsilon)^2} \right) \right), \\ P_{2,n}(\varepsilon) &:= \frac{1}{4} + \frac{P_{3,n}(\varepsilon)}{576}, \end{split}$$

$$\begin{split} P_{3,n}(\varepsilon) &:= \frac{96\sqrt{2\varepsilon}|\lambda_{3,n}|}{(K_{4,n}^{1/4}n^{1/4})} + 48\varepsilon \left(\frac{K_{4,n}}{n}\right)^{1/2} + \frac{32\varepsilon\lambda_{3,n}^2}{(K_{4,n}n)^{1/2}} + \frac{16\sqrt{2}K_{4,n}^{1/4}|\lambda_{3,n}|\varepsilon^{3/2}}{n^{3/4}} + \frac{4\varepsilon^2 K_{4,n}}{n}, \\ e_{2,n}(\varepsilon) &:= \exp\left(\varepsilon^2 \left(\frac{1}{6} + \frac{2P_{2,n}(\varepsilon)}{(1-3\varepsilon)^2}\right)\right), \\ U_{2,2,n}(t) &:= \frac{|t|^5|\lambda_{3,n}|}{6n^{3/2}} + \frac{t^6 K_{4,n}}{24n^2} + \frac{t^6\lambda_{3,n}^2}{36n^2} + \frac{|t|^7 K_{4,n}|\lambda_{3,n}|}{72n^{5/2}} + \frac{t^8 K_{4,n}^2}{576n^3}. \end{split}$$

Note that

$$e_{2,n}(\varepsilon) = e^{\varepsilon^2/6} \times \exp\left(\frac{2\varepsilon^2 P_{2,n}(\varepsilon)}{(1-3\varepsilon)^2}\right) = e_3(\varepsilon) \exp\left(\frac{2\varepsilon^2 P_{3,n}(\varepsilon)}{576(1-3\varepsilon)^2}\right),$$

where $e_3(\varepsilon) := e^{\varepsilon^2/6 + \varepsilon^2/(2(1-3\varepsilon))^2}$.

Lemma 4.6. Under Assumption 4.2, for every $\varepsilon \in (0, 1/3)$ and t such that $|t| \leq \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}$,

$$\left| f_{S_n}(t) - e^{-\frac{t^2}{2}} \left(1 - \frac{it^3 \lambda_{3,n}}{6\sqrt{n}} \right) \right| \le e^{-t^2/2} \left\{ \frac{t^4 K_{4,n}}{8n} \left(\frac{1}{3} + \frac{1}{(1 - 3\varepsilon)^2} \right) + \frac{e_{2,n}(\varepsilon)|t|^6 |\lambda_{3,n}|^2}{72n} + R_{2,n}(t,\varepsilon) \right\}.$$

Proof of Lemma 4.6: This proof is very similar to that of Lemma 4.5. We note that $B_n = \sigma \sqrt{n}$. As before, using two Taylor-Lagrange expansions successively, we can write that for every $\varepsilon \in (0, 1/3)$ and t such that $|t| \leq \sqrt{2\varepsilon n}/K_{4,n}^{1/4}$

$$\log(f_{P_{X_n}}(t/B_n)) = U_{1,n}(t) - \frac{U_{1,n}(t)^2}{2(1+\theta_{2,n}(t)U_{1,n}(t))^2},$$

where

$$U_{1,n}(t) := -\frac{t^2}{2n} - \frac{i\lambda_{3,n}t^3}{6n^{3/2}} + \frac{\theta_{1,n}(t)K_{4,n}t^4}{24n^2},$$

and $\theta_{1,n}(t)$ and $\theta_{2,n}(t)$ are two complex numbers with modulus bounded by 1. Using a third Taylor-Lagrange expansion, we can write that for some complex $\theta_{3,n}(t)$ with modulus bounded by $\exp\left(\frac{K_{4,n}t^4}{24n} + \frac{n|U_{1,n}(t)|^2}{2(1-3\varepsilon)^2}\right)$, the following holds

$$f_{S_n}(t) = e^{-\frac{t^2}{2}} \left(1 - \frac{it^3 \lambda_{3,n}}{6\sqrt{n}} + \frac{t^4 K_{4,n} \theta_{1,n}(t)}{24n} - \frac{nU_{1,n}(t)^2}{2(1 + \theta_{2,n}(t)U_{1,n}(t))^2} + \frac{\theta_{3,n}(t)}{2} \left(-\frac{it^3 \lambda_{3,n}}{6\sqrt{n}} + \frac{t^4 K_{4,n} \theta_{1,n}(t)}{24n} - \frac{nU_{1,n}(t)^2}{2(1 + \theta_{2,n}(t)U_{1,n}(t))^2} \right)^2 \right).$$

Using the triangle inequality and its reverse version plus the condition $|t| \leq \sqrt{2\varepsilon} (n/K_{4,n})^{1/4}$, we obtain

$$\left| f_{S_n}(t) - e^{-t^2/2} \left(1 - \frac{it^3 \lambda_{3,n}}{6\sqrt{n}} \right) \right| \le e^{-t^2/2} \left\{ \frac{t^4 K_{4,n}}{24n} + \frac{nU_{1,n}(t)^2}{2(1-3\varepsilon)^2} + \frac{1}{2} \exp\left(\frac{\varepsilon^2}{6} + \frac{n|U_{1,n}(t)|^2}{2(1-3\varepsilon)^2} \right) \times \left(\frac{|t|^3|\lambda_{3,n}|}{6\sqrt{n}} + \frac{t^4 K_{4,n}}{24n} + \frac{nU_{1,n}(t)^2}{2(1-3\varepsilon)^2} \right)^2 \right\}.$$
(4.51)

We can decompose $nU_{1,n}(t)^2$ as

$$nU_{1,n}(t)^{2} = \frac{t^{4}}{4n} + \underbrace{\frac{|t|^{5}|\lambda_{3,n}|}{6n^{3/2}} + \frac{t^{6}K_{4,n}}{24n^{2}} + \frac{t^{6}\lambda_{3,n}^{2}}{36n^{2}} + \frac{|t|^{7}K_{4,n}|\lambda_{3,n}|}{72n^{5/2}} + \frac{t^{8}K_{4,n}^{2}}{576n^{3}}}_{=U_{2,2,n}(t)}$$

$$\leq \frac{t^{4}P_{2,n}(\varepsilon)}{n}, \qquad (4.52)$$

Combining Equations (4.51) and (4.52) and grouping terms, we conclude that for every $\varepsilon \in (0, 1/3)$ and t such that $|t| \leq \sqrt{2\varepsilon} (n/K)^{1/4}$,

$$\left| f_{S_n}(t) - e^{-\frac{t^2}{2}} \left(1 - \frac{it^3 \lambda_{3,n}}{6\sqrt{n}} \right) \right| \\ \leq e^{-t^2/2} \left\{ \frac{t^4 K_{4,n}}{8n} \left(\frac{1}{3} + \frac{1}{(1-3\varepsilon)^2} \right) + \frac{e_{2,n}(\varepsilon)t^6 \lambda_{3,n}^2}{72n} + \frac{U_{2,2,n}(t)}{2(1-3\varepsilon)^2} \right. \\ \left. + e_{2,n}(\varepsilon) \left(\frac{t^8}{8n^2} \left(\frac{K_{4,n}}{12} + \frac{P_{2,n}(\varepsilon)}{(1-3\varepsilon)^2} \right)^2 + \frac{|t|^7 |\lambda_{3,n}|}{12n^{3/2}} \left(\frac{K_{4,n}}{12} + \frac{P_{2,n}(\varepsilon)}{(1-3\varepsilon)^2} \right) \right) \right\},$$

where $e_{2,n}(\varepsilon) := \exp\left(\varepsilon^2 \left(\frac{1}{6} + \frac{2P_{2,n}(\varepsilon)}{(1-3\varepsilon)^2}\right)\right)$. \Box

4.B.4 Two bounds on incomplete Gamma-like integrals

For every $p \ge 1$, l < m and T > 0, we define J_1 and J_2 by

$$J_1(p, l, m, T) := \frac{1}{T} \int_l^m |\Psi(u/T)| \, u^p e^{-u^2/2} du$$

$$J_2(p, l, m, \widetilde{K}_{3,n}, K_{4,n}, T, n) :=$$
(4.53)

$$\frac{1}{T} \int_{v}^{w} |\Psi(u/T)| u^{p} \exp\left(-\frac{u^{2}}{2} \left(1 - \frac{2\chi_{1}|u|\widetilde{K}_{3,n}}{\sqrt{n}} - \sqrt{\frac{K_{4,n}}{n}}\right)\right) du.$$
(4.54)

We have

$$J_{1}(p, l, m, T) \leq \frac{1.0253}{2\pi} \int_{l}^{m} u^{p-1} e^{-u^{2}/2} du$$

$$\leq \frac{1.0253}{2\pi} \int_{0}^{+\infty} u^{p-1} e^{-u^{2}/2} du$$

$$= \frac{1.0253}{2\pi} \int_{0}^{+\infty} \sqrt{2v}^{p-1} e^{-v} \frac{dv}{\sqrt{2v}} = \frac{1.0253 \times 2^{p/2-2} \Gamma(p/2)}{\pi}.$$
 (4.55)

To obtain (4.55), we resort to the first inequality in (4.9) and the change of variable $v = u^2/2$, and we let $\Gamma(a)$ stand for $\Gamma(a, 0)$.

Exponential decay of the term J_2 .

Using the first inequality in (4.9), we get

$$J_{2}(p,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T/\pi,\widetilde{K}_{3,n},K_{4,n},T,n) \leq \frac{1}{T}\int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{T/\pi}\frac{1.0253\,T}{2\pi u}\,u^{p}\exp\left(-\frac{u^{2}}{2}\left(1-\frac{2\chi_{1}|u|\widetilde{K}_{3,n}}{\sqrt{n}}-u^{2}\sqrt{\frac{K_{4,n}}{n}}\right)\right)du.$$

We now use our choice of T which leads to

$$J_{2}(p,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T/\pi,\widetilde{K}_{3,n},K_{4,n},T,n)$$

$$\leq \frac{1.0253}{2\pi} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{2\sqrt{n}/\widetilde{K}_{3,n}} u^{p-1} \exp\left(-\frac{u^{2}}{2}\left(1-\frac{2\chi_{1}|u|\widetilde{K}_{3,n}}{\sqrt{n}}-u^{2}\sqrt{\frac{K_{4,n}}{n}}\right)\right) du,$$

$$\leq \frac{1.0253}{2\pi} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{2\sqrt{n}/\widetilde{K}_{3,n}} u^{p-1} \exp\left(-\frac{u^{2}}{2}\left(1-4\chi_{1}-\left(\frac{K_{4,n}}{n}\right)^{1/2}\right)\right) du.$$

Note that $1 - 4\chi_1 - (K_{4,n}/n)^{1/2} > 1/4$ when $n > K_{4,n}/(0.75 - 4\chi_1)^2$, i.e., as soon as $n \ge 8 K_{4,n}$. When this is the case, we have

$$\begin{split} J_{2}(p,\sqrt{2\varepsilon}(n/K_{4,n})^{1/4},T/\pi,\widetilde{K}_{3,n},K_{4,n},T,n) \\ &\leq \frac{1.0253}{2\pi} \int_{\sqrt{2\varepsilon}(n/K_{4,n})^{1/4}}^{2\sqrt{n}\widetilde{K}_{3,n}} u^{p-1} \exp\left(-\frac{u^{2}}{8}\right) du \\ &\leq \frac{1.0253}{2\pi} \int_{\varepsilon(n/K_{4,n})^{1/2}/4}^{n/(2\widetilde{K}_{3,n}^{2})} 8^{(p-1)/2} v^{(p-1)/2} e^{-v} \sqrt{\frac{2}{v}} dv \\ &= \frac{1.0253 \times 8^{(p-1)/2} \sqrt{2}}{2\pi} \left(\Gamma\left(p/2,\frac{\varepsilon}{4}\left(\frac{n}{K_{4,n}}\right)^{1/2}\right) - \Gamma\left(p/2,\frac{n}{2\widetilde{K}_{3,n}^{2}}\right) \right) \\ &= O\left(\Gamma\left(p/2,\frac{\varepsilon}{4}\left(\frac{n}{K_{4,n}}\right)^{1/2}\right)^{\frac{p}{2}-1} e^{-\varepsilon(n/K_{4,n})^{1/2/4}} \left(1+O(n^{-1/2})\right) \right. \\ &\left. - \left(\frac{n}{2\widetilde{K}_{3,n}^{2}}\right)^{\frac{p}{2}-1} e^{-n/(2\widetilde{K}_{3,n}^{2})} \left(1+O(n^{-1})\right) \right) \\ &= O(n^{p/4-1/2} e^{-\varepsilon\sqrt{n}/4} \sqrt{K_{4,n}}), \end{split}$$

where we use the change of variable $v = u^2/8$.

4.B.5 Proof of Proposition 4.4

The assumed integrability condition implies that f_Q is absolutely integrable, and therefore we can apply the inversion formula (Ushakov, 2011, Theorem 1.2.6) so that for any $x \in \mathbb{R}$,

$$q(x) = \int_{-\infty}^{+\infty} r(x,t)dt.$$

where $r(x,t) := \frac{1}{2\pi} e^{-itx} f_P(t)$. Note that r is infinitely differentiable with respect to x, and that

$$\left|\frac{\partial r(x,t)}{\partial x^{p-1}}\right| = \left|\frac{1}{2\pi}(-it)^{p-1}e^{-itx}f_P(t)\right| = \frac{1}{2\pi}|t|^{p-1}|f_P(t)|,$$

which is integrable with respect to t, by assumption. This concludes the proof that q is (p-1) times differentiable, as r is measurable.

Chapter 5

On the construction of nonasymptotic confidence intervals in linear models

This fifth chapter uses the result of the fourth one to construct nonasymptotic confidence intervals in linear models. It is again joint work with Alexis Derumigny and Yannick Guyonvarch.

Abstract We are concerned with constructing nonasymptotic confidence intervals (CIs) for linear functionals of the vector of coefficients in a linear regression model, i.e., CIs that have a probability of containing the true parameter at least equal to their nominal level for any sample size. In the linear regression setting with exogenous regressors, we propose a novel closed-form CI that is tightly connected to the one based on the standard *t*-statistic. Unlike the latter, which has only asymptotic theoretical validity, ours is valid for any sample size, uniformly over a large class of distributions defined by moment restrictions only (nonparametric). In particular, we allow for heteroskedasticity but require the influence function of the OLS estimator to have bounded kurtosis. Furthermore, the length of our CI converges in probability to the one of the standard CI based on the *t*-statistic, so that our inferential procedure is asymptotically exact. We investigate the practical performance of our method in a simulation study. We extend our theoretical results to linear models with endogeneity: with one endogenous regressor, one instrument, and additional exogenous covariates, we propose a modification of the asymptotic Anderson-Rubin confidence set (CS) that is valid for any sample size under analogous conditions to the exogenous case. Our work builds upon a long-standing statistics literature that was crucially influenced by Esseen (1945) and Cramer (1962).

5.1 Introduction

A large part of applied works in economics uses linear regression models. In general, interest does not lie in the entire vector of parameters but in some individual coefficients, typically the coefficient of a treatment status indicator. In a vast majority of applications, confidence intervals (CIs) for such coefficients are based on the asymptotic normality of the corresponding t-statistic, which requires minimal moment conditions on regressors and residuals. Consequently, their theoretical guarantees hold only at the limit, when the number of observations tends to infinity.

From a theoretical viewpoint, it has been established that CIs based on the *t*-statistic have poor properties in finite samples without additional assumptions (see Bahadur and Savage (1956), Dufour (1997), or Bertanha and Moreira (2016)). More precisely, the probability that they contain the true parameter (i.e. their coverage) falls below their stated nominal level. To circumvent this failure, one can impose that the residuals of the model (I) are independent of the regressors, and (N) follow a Gaussian distribution $\mathcal{N}(0, \sigma^2)$. However, those conditions can often be deemed too restrictive. In particular, the first one rules out heteroskedasticity while the second cannot encompass heavy-tailed or skewed distributions although various economic variables display such behavior.¹

Those limitations have prompted a whole strand of research in econometrics and statistics which aims at constructing nonasymptotic CIs without relying on (I) and (N). Nonasymptotic means that, for any sample size, the CI should contain the true parameter with a probability equal or larger to a user-specified threshold $1 - \alpha$, with $\alpha \in (0, 1)$. We evoke some of these approaches below. Here, we briefly highlight the main challenges faced by nonasymptotic CIs.

The first one relates to their lengths. A very coarse method to address the lack of coverage would be to increase the length of the usual CI based on asymptotic normality, multiplying it by, say, two. By doing so, we would mechanically be less worried about potential departures from their nominal level in finite samples. However, the resulting CI would suffer from being conservative asymptotically: at the limit, it contains the true parameter with a probability that exceeds the desired level $1 - \alpha$. This is unappealing: by setting a nominal level to $1 - \alpha$, we are ready and willing to accept that our CIs miss the parameter with probability α in exchange for an increased precision through smaller CIs' lengths. A second challenge when constructing CIs relates to uniformity. Since we do not know the true distribution of the data, we want CIs which attain a minimal coverage whatever this distribution is, provided it belongs to a given class of distributions that constitutes our statistical model.

The two issues interact. Ideally, we are looking for a CI whose coverage is exactly equal to $1 - \alpha$ for any sample size (hence said to be exact instead of conservative),² and for any distribution within a class as large as possible. The impossibility results presented in Bahadur and Savage (1956), Dufour (1997), or Bertanha and Moreira (2016) extend beyond the *t*-statistic and notably imply that such a general objective cannot be achieved.

¹Here, "heavy-tailed" refers to distributions whose moments are not finite from a certain order contrary to Gaussian ones. For aggregate macroeconomic variables, see for instance Acemoglu et al. (2017) and Ascari et al. (2015); for individual microeconomic variables, an example is the modelling of wage distributions by a Pareto distribution, whose number of finite moments depends on its shape parameter.

²The term "exact" is sometimes also used to refer to finite-sample properties. In this paper, we use the words "nonasymptotic" for the latter and "exact" in opposition to "conservative".

More precisely, a trade-off arises between how sharp a CI's properties can be (typically being exact, not conservative) and the richness of the class of distributions the CI applies to. The coverage generally depends on some specific characteristics of the data-generating process. This is why achieving a coverage exactly equal to some given threshold in finite samples entails severe restrictions on the class of distributions we consider (*cf.* assumptions (I) and (N) above or Example 5.1 that illustrates our discussion in the simpler example of conducting inference for an expectation). In that sense, requiring nonasymptotic exactness is very demanding.

On the other hand, it is more reasonable to ask for asymptotic exactness. A formal definition is given in Section 5.2. As for now, consider that a CI is said to be asymptotically exact if, whatever the distribution of the data inside a given class, the limit when the sample size tends to infinity of its coverage is equal to $1 - \alpha$.³ Asymptotic exactness can be used as an optimality criterion to choose which inferential procedures are the best among a set of admissible CIs. The usual CI relying on the asymptotic normality of the *t*-statistic is asymptotically exact. As a result, a "good" CI should behave like the one based on the *t*-statistic at the limit.

In this paper, our objective is to construct nonasymptotic CIs in linear regression models, having guaranteed coverage for any sample size, and that are asymptotically exact without relying on the restrictive assumptions of independent and normal residuals.

We first consider linear models without endogenous regressors and seek to conduct inference on scalar parameters of the form $u'\beta_0$, where u is a known vector and β_0 corresponds to the coefficients of the linear regression under investigation.⁴ Leveraging results for self-normalized sums and bounds on Edgeworth expansions, we derive CIs for such parameters under moment conditions only. In particular, we impose two moments restrictions on the joint distribution of regressors and residuals: a bounded fourth moment for the residuals times regressors and a bounded kurtosis for the influence function of the OLS estimator. Compared to the usual CI based on the asymptotic normality of the t-statistic, these conditions can be deemed the main additional ones. Indeed, the other restrictions we rely on essentially strengthen the usual moment conditions of linear models, putting bounds that are somewhat unavoidable to achieve nonasymptotic properties in nonparametric statistical models. We also present an improved version of our CI under additional regularity conditions, which are closely connected to the residuals and regressors having a continuous distribution with respect to Lebesgue's measure. To the best of our knowledge, our procedure is the first to meet the following four criteria at the same time in linear regression models without endogenous regressors: (i) nonasymptotic validity, (ii)allowing for heteroskedasticity, (*iii*) asymptotic exactness, and (*iv*) having a closed form expression.

In addition to its theoretical properties, our procedure can prove interesting in practice. The closed-form expression of our CI makes it immediate to compute. Nonetheless, our method requires knowledge of bounds. Such types of inputs are common to build nonasymptotic procedures. Practitioners can use external information to state sensible values for those bounds. When lacking such a priori knowledge, we propose a plug-in strategy to replace those unknown bounds and discuss its relevance. Some simulation studies reveal the possible usefulness of our procedure in real applications. In moderate

³The important part is "equal", in opposition to "larger or equal" for conservative CIs.

⁴We note that conducting inference on individual coefficients of the linear model falls in that category.

(thousands of observations) to small (hundreds) samples, although broader than the classical *t*-statistic's one in order to ensure finite-sample coverage, our CI appears sufficiently precise to be informative about the parameter of interest.

We finally extend our theoretical analysis to cases with one endogenous regressor, one instrument, and exogenous covariates. This configuration does not exhaust linear models with endogeneity. Nonetheless, according to the recent review of Andrews et al. (2019), it is the leading case. In this setting, we construct a modified version of the Anderson and Rubin (1949)'s confidence set (CS).⁵ As in the exogenous case, under similar moment conditions, our CS has guaranteed coverage for any sample size and simultaneously is asymptotically exact. Moreover, similar to Anderson and Rubin's CS, it remains valid even in the weak instrumental variable (IV) case: when the instrument is uncorrelated with the endogenous variable.

The rest of the paper is organized as follows. A literature review completes this introduction followed by the notation used in the paper. Section 5.2 defines several characteristics of confidence sets that enable to gauge their performance. It also discusses in detail how to construct CIs for an expectation to illustrate the major trade-offs and strategies in a simple case. Our CI in exogenous linear models is defined and shown to be nonasymptotically valid as well as asymptotically exact in Section 5.3. In Section 5.4, we discuss its practical implementation. Some simulations illustrate the behavior of our CI in comparison to the one based on t-statistic's asymptotic normality in Section 5.5. We present our theoretical findings on IV models in Section 5.6. Appendices 5.A and 5.B present the proofs of our main results. The additional lemmas used are presented in Appendix 5.C.

Literature The question of conducting nonasymptotic inference has a long history in statistics and econometrics. Inference on a (multivariate) mean is of course the problem that has received the most attention. In the model $Y = \mu + \epsilon$ with $\mu \in \mathbb{R}, \epsilon \sim \mathcal{N}(0, \sigma^2)$, and $\sigma > 0$, Student (1908) was the first to show the t-statistic based on *i.i.d.* observations follows exactly a Student distribution in finite samples. Berry (1941) and Esseen (1942) are two other early seminal contributions which give a nonasymptotic bound on the distance between the distribution of a sample mean based on i.i.d. observations and a Gaussian distribution, a result called the Berry-Esseen inequality. Building on this inequality, Hall and Jing (1995) and Romano and Wolf (2000) proposed nonasymptotic CIs for a mean requiring only light nonparametric constraints on the data distribution. Those CIs have the additional nice property of being asymptotically exact but remain difficult to use in practice: they depend on unspecified universal constants that could prove difficult to estimate/compute, or simply be very large. Romano and Wolf (2000) is also one of the few existing contributions discussing in a clear and exhaustive manner the challenge and importance of conducting nonasymptotic inference which seeks to be asymptotically exact. The Berry-Esseen inequality was extended in several directions from the 1990s on. First, Bentkus and Götze (1996) and Bentkus et al. (1996) proposed a version of this inequality for the t-statistic. Shao (2005) improved on the last two contributions by giving explicit numerical constants in the Berry-Esseen inequality for the t-statistic. While Shao (2005)'s result directly allows to construct a nonasymptotic confidence interval (also asymptotically exact) for a mean modulo some higher-moments' bounds, the numerical constants are still too large to obtain informative CIs in practice (see also Pinelis (2011a)

⁵This procedure does not always produce an interval, hence the term "confidence set".

for a discussion on this issue and some further results). The Berry-Esseen inequality was more recently adapted to multivariate frameworks (Bentkus (2003), Raič (2019)). Similar to the *t*-statistic case, those elegant theoretical contributions display too crude numerical constants, preventing their application. However, as noted in Raič (2019), there seems to remain a lot of room to obtain smaller numerical constants in the multivariate setup. If the data distribution is symmetric about its expectation, it is possible to conduct nonasymptotic inference on the expectation of the data distribution without imposing more moments. This was first observed by Efron (1969), refined recently in Bentkus and Dzindzalieta (2015), and was extended to multivariate setups by Pinelis (1994). This approach has one major drawback: it is not asymptotically exact. Bertail et al. (2008) show that the method can be adapted (under additional moment constraints) when the data does not display symmetry.

Nonasymptotic inference in regression models has emerged more recently. In linear models without endogeneity, Gossner and Schlag (2013) propose two nonasymptotic inference methods, valid under boundedness of the outcome variable. This allows in particular for heteroskedastic residuals. There are several limitations though, either theoretical or practical: the first method is not exact asymptotically and neither the first nor the second method yield closed-form CIs or CSs. In the statistics literature, Kuchibhotla et al. (2019) and Kuchibhotla et al. (2020) have recently proposed methods to conduct valid inference in finite samples on subsets of coefficients in a possibly highdimensional linear regression framework. Kuchibhotla et al. (2020) even provide guarantees for inference after a model-selection step. While these results are very general and are shown under mild moment assumptions on the data distribution, they do not lead to asymptotically exact inference on individual parameters of the linear regression for instance. Still in linear models without endogeneity, DiCiccio and Romano (2017) advocate the use of permutation-based inference. When the regressors are jointly independent from the residual, they show that permutation-based inference leads to exact nonasymptotic inference on the full vector of coefficients (except the constant). Their condition therefore rules out heteroskedastic residuals. When one is interested in a subset of the coefficients, the authors suggest an alternative procedure which is nonasymptotically exact and requires a quite stringent condition: the regressors associated with the coefficients of interest must be independent from the residual and the remaining regressors. This places strong restrictions on possible heteroskedasticity as well as unnatural assumptions on the link between regressors. In a general M-estimation framework, Schreuder et al. (2020) derive nonasymptotic confidence intervals with closed-form expression on individual components of the coefficients' vector. They only require nonparametric restrictions on the data distribution. Their result is much stronger than what is usually required in the following sense: with probability at least $1 - \alpha$, their CI of target level $1 - \alpha$ contains the true parameter for every sample size large enough at the same time. The main drawback of their result is a lack of asymptotic exactness. Chernozhukov et al. (2009) tackle the issue of conducting inference in a generic quantile regression framework allowing for endogenous regressors. They propose a simulation-based confidence set for the full vector of coefficient that is exact nonasymptotically. The main restriction they impose is independence between residuals and exogenous regressors plus excluded instruments. They thus do not allow for heteroskedasticity. Moreover, their CSs do not have a closed-form formula and are computationally demanding. Besides, when one is interested in individual coefficients, their tools produce CSs that are not asymptotically exact.

Our work is in addition closely connected to the literature on weak instrumental variables

(weak IVs). Anderson and Rubin (1949) was an early contribution to this literature. In this work, the authors introduce a method (that we call the AR confidence sets) to conduct inference in linear IV models with as many instruments as endogenous regressors. Their method does not require instruments and endogenous variables to be correlated. The AR confidence sets are obtained by inverting a collection of Wald tests. The test-inversion rationale was extended in numerous ways throughout the years: for instance, Kleibergen (2002) replaced Wald tests with score tests and Moreira (2003) introduced likelihood ratio tests. The theoretical properties of those methods have been thoroughly investigated (i) asymptotically (see Andrews and Mikusheva (2016) or Andrews and Guggenberger (2019)), and (ii) nonasymptotically under the assumption that residuals are normally distributed with known variance (Andrews et al. (2019)). Tuvaandorj (2020) has recently proposed to invert permutation-based versions of the tests introduced in Anderson and Rubin (1949), Kleibergen (2002) or Moreira (2003): his method is valid nonasymptotically as soon as exogenous regressors and excluded instruments are independent from the second stage residuals and valid asymptotically under much weaker conditions. To the best of our knowledge, all existing nonasymptotic results in the weak IV setup rely either on the assumption that residuals are normally distributed or that they are independent from exogenous covariates and excluded instruments.

Notation For a random variable D, we denote by P^D its distribution and by support(D)or support (\mathbf{P}^D) its support. Similarly $\mathbf{P}^{D,U}$ denotes the joint distribution of a pair of random variables (D, U). For a parameter θ , associated to a given statistical model, P_{θ} denotes a distribution indexed by the parameter. $P_{D,U} = P_D \otimes P_U$ means that D and U are independent. For any set $\mathcal{D}, \mathcal{P}(\mathcal{D})$ denotes the set of all probability distribution supported on \mathcal{D} . For any real vector $u = (u_1, \ldots, u_d), ||u|| := (u_1^2 + \ldots + u_d^2)^{1/2}$ denotes its Euclidean or ℓ_2 -norm. For matrices, we consider the operator norm induced by the Euclidean norm: for any real matrix M, ||M|| denotes its spectral norm, namely the square root of the largest eigenvalue of M'M, with M' the transpose of M. We also denote by $\lambda_{\min}(M)$ the smallest eigenvalue of M. $vec(\cdot)$ denotes the vectorization of a matrix, that is, if M is a $m \times n$ matrix, vec(M) is the $mn \times 1$ column vector obtained by stacking the columns of the matrix M on top of one another. For any distribution P and real number $\tau \in (0,1), q_{\rm P}(\tau)$ denotes the quantile at order τ of the distribution. Since we remain in an *i.i.d.* set-up throughout the article, we sometimes drop the subscript i of random variables to lighten notations. In other words, if D_1, \ldots, D_n are *n i.i.d.* random variables, D without subscript denotes a generic random variable with the same distribution.

5.2 Quality measures for confidence sets

We start by formally defining several attributes of confidence sets that enable to characterize their quality. We do so in a general framework, which encompasses in particular linear models with and without endogeneity. In the next sections, we reduce to those two settings.

As for now, let us assume we observe $n \in \mathbb{N}^*$ independent and identically distributed (i.i.d.) data points $(D_i)_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}^D$ which are vectors of dimension $d \in \mathbb{N}^*$. We denote by $\mathcal{D} \subseteq \mathbb{R}^d$ the support of the distribution of D. Throughout the article, we assume that the distribution of the data belongs to a statistical model whose parameter can be divided into a finite-dimensional part and the remainder:

$$\mathbf{P}_{D} \in \left\{ \mathbf{P}_{\theta} \in \mathcal{P}(\mathcal{D}), \theta = (\theta_{1}, \theta_{2}) \in \Theta := \mathbb{R}^{p} \times \Theta_{2} \right\},\$$

where Θ_2 is a topological space and $p \in \mathbb{N}^*$. The expression just represents that typical models considered by econometricians are semi-parametric, in the sense that the distribution of the data is characterized by both a finite-dimensional parameter, which is the one of interest, and an infinite-dimensional "nuisance" parameter, such as some set of probability distributions. For instance, in linear regression models, Θ_2 is a nonparametric set of joint distributions for regressors and residuals.

We seek to construct a confidence set (CS) for a scalar functional $T(\cdot)$ of θ_1 . Formally, a CS for $T(\theta_1)$ is a measurable map from \mathcal{D}^n to the set of closed subsets of \mathbb{R} – *e.g.*, closed intervals or unions of closed intervals are closed subsets of \mathbb{R} . When evaluated at $(D_i)_{i=1}^n$, a CS is a random closed subset of \mathbb{R} that depends on the data. Confidence intervals are a special type of CSs that need to be intervals. A confidence set comes with a desired confidence or nominal level $1 - \alpha$, for some $\alpha \in (0, 1)$. Indeed, its main objective is basically to contain its target $T(\theta_1)$ with probability at least $1 - \alpha$. Hereafter, the notation $CS(1 - \alpha, n)$ stands for a generic confidence set to denote the dependence in $(D_i)_{i=1}^n$ and α .

Let Θ be a fixed subset of Θ . Several criteria exist to assess the quality of $CS(1-\alpha, n)$. From an asymptotic point of view, $CS(1-\alpha, n)$ is said to be asymptotically exact pointwise over Θ if

$$\forall \theta \in \widetilde{\Theta}, \lim_{n \to +\infty} \mathbb{P}_{\mathbb{P}_{\theta}^{\otimes n}} \Big(\mathrm{CS}(1-\alpha, n) \ni T(\theta_1) \Big) = 1 - \alpha.$$
(5.1)

A stronger asymptotic criterion exists: $CS(1 - \alpha, n)$ is said to be asymptotically exact uniformly over $\tilde{\Theta}$ if

$$\lim_{n \to +\infty} \sup_{\theta \in \widetilde{\Theta}} \left| \mathbb{P}_{\mathbf{P}_{\theta}^{\otimes n}} \left(\mathrm{CS}(1-\alpha, n) \ni T(\theta_1) \right) - (1-\alpha) \right| = 0.$$
(5.2)

If $\liminf_{n \to +\infty} \inf_{\theta \in \widetilde{\Theta}} \mathbb{P}_{\mathbb{P}_{\theta}^{\otimes n}} (CS(1 - \alpha, n) \ni T(\theta_1)) \geq 1 - \alpha$ the CS is asymptotically conservative uniformly over $\widetilde{\Theta}$.

From a nonasymptotic perspective, we say that $CS(1 - \alpha, n)$ is nonasymptotically conservative over $\widetilde{\Theta}$ if

$$\forall n \ge 1, \inf_{\theta \in \widetilde{\Theta}} \mathbb{P}_{\mathsf{P}_{\theta}^{\otimes n}} \Big(\mathrm{CS}(1-\alpha, n) \ni T(\theta_1) \Big) \ge 1-\alpha.$$
(5.3)

This property evolves into nonasymptotic exactness over $\tilde{\Theta}$ if (5.3) holds and in addition for any $n \geq 1$, $\sup_{\theta \in \tilde{\Theta}} \mathbb{P}_{\mathbf{P}_{\theta}^{\otimes n}} \left(\mathrm{CS}(1-\alpha, n) \ni T(\theta_1) \right) = 1-\alpha$.

The quality of CSs can thus be evaluated along several dimensions. For a given nominal level, the opposition between exact and conservative relates to the actual precision of CSs. Another distinction contrasts nonasymptotic and asymptotic results. In the latter case, we can further distinguish pointwise from uniform properties.⁶ Asymptotic uniform validity relates to the notion of "honesty." Compared to pointwise guarantees, it can be argued as more reliable regarding CSs' finite-sample performance in practice (Armstrong and Kolesár, 2020). In this article, we follow another road to finite-sample validity through direct nonasymptotic results.

Property (5.1) is implicitly what econometricians have in mind when they rely on the asymptotic normality of an estimator to conduct inference. This property is usually

⁶Remark that there is no such distinction in nonasymptotic results: for a fixed n, the inf over Θ in (5.3) is equivalent to a universal quantifier $\forall \theta \in \Theta$.

achievable on the whole of Θ . As illustrated in Kasy (2019) for instance, in models where (5.1) holds, it is often straightforward to define a (large) subset $\tilde{\Theta}$ on which (5.2) is verified. As regards finite-sample inference, property (5.3) has been shown to apply in a large family of models. These results can be found predominantly in the theoretical statistics literature (see Schreuder et al. (2020) for a recent illustration). These are powerful findings which however yield CSs that are in general asymptotically conservative, uniformly and even pointwise.⁷ Finally, the strongest notion, namely nonasymptotic exactness, can be obtained in a wide class of models at the cost of placing fairly strong restrictions on P_{θ} , for instance by imposing that P_{θ} belongs to a parametric family. We illustrate this discussion in Example 5.1 and Section 5.3.

Example 5.1 (Inference on an expectation). We are interested in conducting inference on $T(\theta_1) = \theta_1 := \mathbb{E}[D] \in \mathbb{R}$. In this example, $\theta_2 := P_D$ and $\Theta_2 := \{P \in \mathcal{P}(\mathbb{R}) : \int u^2 dP < +\infty\}.$

Letting $\hat{\sigma}_n^2 := (n-1)^{-1} \sum_{i=1}^n (D_i - \overline{D}_n)^2$, with $\overline{D}_n := n^{-1} \sum_{i=1}^n D_i$, we know by the Central Limit Theorem (CLT) and Slutsky's lemma that

$$CS_1(1-\alpha,n) := \left[\overline{D}_n \pm \frac{q_{\mathcal{N}(0,1)}(1-\alpha/2)\sqrt{\widehat{\sigma}_n^2}}{\sqrt{n}}\right]$$

is asymptotically exact pointwise over Θ .

Following Kasy (2019) (in particular Proposition 1), $CS_1(1-\alpha, n)$ is also asymptotically exact uniformly over $\tilde{\Theta}_1 := \{\theta \in \Theta : \mathbb{V}_{P_{\theta}}(D) \ge m, \mathbb{E}_{P_{\theta}}[D^4] \le M\}$ with $m > 0, M < +\infty$. Hence, the uniform property holds on a large subset of Θ insofar as $\tilde{\Theta}_1$ is nonparametric. On the other hand, results by Bahadur and Savage (1956) and Dufour (1997) imply that $CS_1(1-\alpha, n)$ cannot be nonasymptotically conservative over Θ nor asymptotically exact uniformly over Θ .

This negative result is somewhat counterbalanced by the well-known fact that, up to using the quantiles of a Student instead of a standard Normal distribution, $CS_1(1-\alpha, n)$ is nonasymptotically exact provided P_D is a Gaussian distribution. Denoting \mathcal{T}_{n-1} a Student distribution with n-1 degrees of freedom,

$$CS_2(1-\alpha,n) := \left[\overline{D}_n \pm \frac{q_{\tau_{n-1}}(1-\alpha/2)\sqrt{\widehat{\sigma}_n^2}}{\sqrt{n}}\right]$$

is nonasymptotically exact over $\widetilde{\Theta}_2 := \left\{ \theta \in \Theta : \mathbb{P}^{\theta} = \mathcal{N}(\theta_1, \sigma^2), \theta_1 \in \mathbb{R}, \sigma^2 \in \mathbb{R}^*_+ \right\}$. It is also asymptotically exact pointwise over Θ , as $\mathrm{CS}_1(1 - \alpha, n)$. Remark that in this case, nonasymptotic exactness is obtained without relying on upper and lower bounds on moments.

Another way to nonasymptotic guarantees dispenses with parametric restrictions at the expense of known bounds on some moments (or support) of \mathbb{P}^D . Let $\widetilde{\Theta}_3 := \{\theta \in \Theta : \mathbb{V}_{P_\theta}(D) \leq M\}$, for some $M < +\infty$. An alternative CS for θ_1 is

$$CS_3(1-\alpha,n) := \left[\overline{D}_n \pm \sqrt{\frac{M}{\alpha n}}\right].$$

⁷By saying those CSs are in general asymptotically conservative, even pointwise, we refer to the fact that they are often such that in a subset $\tilde{\Theta}$ of Θ , $\forall \theta \in \tilde{\Theta}$, $\lim_{n \to +\infty} \mathbb{P}_{\mathbb{P}^{\otimes n}} \left(\mathrm{CS}(1-\alpha, n) \ni T(\theta_1) \right) > 1-\alpha$.

Using the Bienaymé-Chebyshev inequality, it can be shown that $\operatorname{CS}_3(1-\alpha, n)$ is nonasymptotically conservative over $\widetilde{\Theta}_3$. $\operatorname{CS}_3(1-\alpha, n)$ requires a known upper bound on the variance. On the other hand, compared to $\widetilde{\Theta}_2$, $\widetilde{\Theta}_3$ can be deemed a large subset of Θ for it is nonparametric. However, this CS has one major drawback: it is asymptotically conservative, even pointwise, over $\widetilde{\Theta}_3$. To see this, one can simply remark that for every $\alpha \in (0,1)$, $1/\sqrt{\alpha} > q_{\mathcal{N}(0,1)}(1-\alpha/2)$. This implies that $1/\sqrt{\alpha} = q_{\mathcal{N}(0,1)}(1-\beta/2)$ for some $\beta < \alpha$. Thus, for any $\theta \in \widetilde{\Theta}_3$, the probability that $\operatorname{CS}_3(1-\alpha, n)$ contains θ_1 has a limit at least $1-\beta > 1-\alpha$ when n goes to infinity.

 $\text{CS}_3(1-\alpha,n)$ is a basic instance of a CS constructed thanks to a concentration inequality.⁸ There exist many different concentration inequalities (see Boucheron et al. (2013) for an in-depth exposition) relying on alternative restrictions $\tilde{\Theta}$ on the distribution of the data, each yielding a CS that is nonasymptotically conservative over the relevant $\tilde{\Theta}$. For another example, thanks to Hoeffding inequality,

$$CS_4(1-\alpha, n) := \left[\overline{D}_n \pm \frac{(b-a)}{2} \frac{\sqrt{2\log(2/\alpha)}}{\sqrt{n}}\right]$$

is nonasymptotically conservative over $\widetilde{\Theta}_4 := \left\{ \theta \in \Theta : \operatorname{support}(\mathrm{P}^{\theta}) \subseteq [a, b] \right\}$ for some $-\infty < a < b < +\infty$. In general, CSs based on concentration inequalities display the same suboptimal asymptotic behavior as $\operatorname{CS}_3(1-\alpha, n)$: for every $\alpha \in (0,1), \sqrt{2\log(2/\alpha)} > q_{\mathcal{N}(0,1)}(1-\alpha/2)$ in the case of $\operatorname{CS}_4(1-\alpha, n)$.

In the simple case of a scalar expectation, Example 5.1 highlights the difficulty to construct accurate CSs with nonasymptotic guarantees on large subset of the parameter space Θ . As $CS_2(1 - \alpha, n)$ illustrates, nonasymptotic exactness can be achieved under parametric restrictions. In contrast, provided known bounds on some moments, concentration inequalities yield nonasymptotic conservativeness on large nonparametric subsets $\tilde{\Theta}$ but the resulting CSs are less precise: they are conservative even asymptotically. More generally, Example 5.1 illustrates the fact that it is challenging to build a CS satisfying the following three criteria at the same time: (*i*) weak restrictions on P^{θ} (in particular allowing for semi-parametric models), (*ii*) nonasymptotic conservativeness over Θ or a large subset of Θ , (*iii*) asymptotic exactness uniformly or pointwise.

5.3 Linear regression without endogeneity

5.3.1 Model and standard asymptotic inference

In this section, we present our confidence interval and its properties in the setting of a linear regression model without endogeneity: the residual of the model of interest is orthogonal with regressors. The following assumptions explicit the statistical model and present the conditions we rely on to construct our CI.

We observe *n* observations $(X_1, Y_1), \ldots, (X_n, Y_n)$ of a *p*-dimensional explanatory variable X and an outcome real random variable Y that satisfy a basic linear regression model as formalized in Assumption 5.1.

⁸By concentration inequalities, we mean results of the form $\mathbb{P}_{\mathbb{P}_{\theta}^{\otimes n}}\left(\left|\overline{D}_{n}-\mathbb{E}[D]\right| > f(n,\alpha)\right) < 1-\alpha, \forall \alpha \in (0,1)$ and for some positive-valued function $f(\cdot,\cdot)$.

Assumption 5.1 (Linear model). $(X_i, Y_i)_{i=1}^n \stackrel{i.i.d.}{\sim} P_{X,Y}$, where X is a random vector of dimension p and Y a real random variable such that, for some random variable ε and vector β_0 :

$$Y = X'\beta_0 + \varepsilon, \ \mathbb{E}[X\varepsilon] = 0, \ \mathbb{E}[||X||^2] < +\infty, \ \mathbb{E}[||X\varepsilon||^2] < +\infty, \ \lambda_{\min}(\mathbb{E}(XX')) > 0 \quad (5.4)$$

The parameter set of the associated statistical model is $\Theta := \{\theta = (\beta_0, P_{X,\varepsilon}) \in \mathbb{R}^p \times \mathcal{P}_{X,\varepsilon}\}$ with $\mathcal{P}_{X,\varepsilon} :=$

$$\Big\{ \mathbf{P}^{X,\varepsilon} \in \mathcal{P}(\mathbb{R}^{p+1}) : \mathbb{E}[X\varepsilon] = 0, \ \mathbb{E}[\|X\|^2] < +\infty, \ \mathbb{E}[\|X\varepsilon\|^2] < +\infty, \ \lambda_{\min}(\mathbb{E}(XX')) > 0 \Big\}.$$

In this model, P^{θ} denotes a distribution of (X, Y) indexed by θ . In what follows, we consider several subsets of Θ characterized by additional restrictions that enable to build interesting confidence sets.

Assumption 5.1 sets a basic linear regression model. Θ is indeed the largest parameter set compatible with usual economic assumptions and minimal statistical conditions. The condition $\mathbb{E}[X\varepsilon] = 0$ imposes (weak) exogeneity of covariates. It corresponds to the orthogonality condition of the linear projection of Y on X. It is implied by the (strong) exogeneity assumption $\mathbb{E}[\varepsilon|X] = 0$ but is more general as the conditional expectation of Y given X need not be linear. The other moments conditions allow for heteroskedasticity while ensuring the asymptotic normality of the Ordinary Least Squares (OLS) estimator of β_0 :

$$\widehat{\beta} := \left(\frac{1}{n} \sum_{i=1}^n X_i X_i'\right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n X_i Y_i\right).$$

Remark that the condition $\lambda_{\min}(\mathbb{E}(XX')) > 0$ is equivalent to the invertibility of the matrix $\mathbb{E}[XX']$. Thus, under Assumption 5.1, $\hat{\beta}$ is well-defined with probability approaching one as the sample size n goes to infinity.

For a given vector u of \mathbb{R}^p , our goal is to build a confidence interval for a linear functional of the form $u'\beta_0$. It encompasses CIs for each individual component of β_0 (taking for u the canonical vectors) but also differences of coefficients that appear when investigating the relative impact of two covariates. We consider henceforth an arbitrary vector $u \in \mathbb{R}^p \setminus \{0_{\mathbb{R}^p}\}$.

As mentioned in the introduction, the standard way to proceed is to construct a CI centered at the estimator $u'\hat{\beta}$ relying on the asymptotic normality of $\hat{\beta}$:

$$\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha,n) := \left[u'\widehat{\beta} \pm \frac{q_{\mathcal{N}(0,1)}(1-\alpha/2)}{\sqrt{n}} \sqrt{u'\widehat{V}u} \right]$$
(5.5)

where

$$\widehat{V} := \left(\frac{1}{n}\sum_{i=1}^{n}X_{i}X_{i}'\right)^{-1} \left(\frac{1}{n}\sum_{i=1}^{n}X_{i}X_{i}'\widehat{\varepsilon}_{i}^{2}\right) \left(\frac{1}{n}\sum_{i=1}^{n}X_{i}X_{i}'\right)^{-1}$$

is the standard estimator of the asymptotic variance of $\hat{\beta}$, $V := \mathbb{E}[XX']^{-1}\mathbb{E}[XX'\varepsilon^2]\mathbb{E}[XX']^{-1}$, with $\hat{\varepsilon}_i := Y_i - X'_i \hat{\beta}$ the estimated residual for the *i*-th observation.

The pros and cons of $CI_u^{Asymp}(1-\alpha, n)$ are well-understood and very close to those of $CS_1(1-\alpha, n)$ in Example 5.1. Its good properties are mainly asymptotic. They require

additional moment conditions to ensure that \hat{V} is consistent for the asymptotic variance of $\hat{\beta}$:

$$\Theta^{\operatorname{Asymp}} := \left\{ \theta \in \Theta : \mathbb{E}[\|X\|^4] < +\infty \right\}.$$

By applications of the Law of Large Numbers, the CLT, and Slutsky's lemma, $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha, n)$ is known to be asymptotically exact pointwise over $\Theta^{\operatorname{Asymp}}$.

Besides, following again Kasy (2019), $CI_u^{Asymp}(1-\alpha, n)$ can be strengthened to be asymptotically exact uniformly over

$$\overline{\Theta}^{\text{Asymp}} := \left\{ \theta \in \Theta : \lambda_{\min} \left(\mathbb{E}[XX'] \right) \ge m, \mathbb{E}[||X||^4] \le M, m \le \mathbb{E}[||X\epsilon||^4] \le M \right\},\$$

with $0 < m \leq M < +\infty$.

Furthermore, similar to $CS_2(1 - \alpha, n)$ in Example 5.1, provided \hat{V} is replaced with the (homoscedastic) estimator

$$\frac{1}{n-p}\sum_{i=1}^{n}\widehat{\varepsilon}_{i}^{2}\left(\frac{1}{n}\sum_{i=1}^{n}X_{i}X_{i}'\right)^{-1}$$

and using the $(1 - \alpha/2)$ -quantile of a Student with n - p degrees of freedom instead of a $\mathcal{N}(0,1)$, $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha,n)$ becomes nonasymptotically exact over

$$\Theta^{\mathcal{N}} := \Big\{ \theta \in \Theta^{\operatorname{Asymp}} : \operatorname{P}_{X,\varepsilon} = \operatorname{P}_X \otimes \operatorname{P}_{\varepsilon}, \operatorname{P}_{\varepsilon} = \mathcal{N}(0,\sigma^2), \sigma^2 > 0 \Big\}.$$

However, the restrictions on P_{θ} to achieve such finite-sample properties are often considered as too restrictive in practice: Gaussian residuals impede skewed or heavier tails shocks and independence between ε and X rules out heteroskedasticity. In what follows, we propose CIs with finite-sample guarantees (nonasymptotic conservativeness) and some form of asymptotic efficiency (asymptotic pointwise exactness) without relying on such independence or parametric assumptions.

5.3.2 Our confidence interval

We introduce the random variable

$$\xi := u' \mathbb{E}[XX']^{-1} X \varepsilon$$

that comes from the linearization of $\hat{\beta}$ used in the construction of our nonasymptotic CIs. The nonasymptotic guarantees on those CIs hold on a subset of Θ described in Assumption 5.2.

Assumption 5.2 (Bounds on DGP). The parameter $\theta = (\beta_0, P^{X,\varepsilon})$ belongs to Θ and there exist positive constants λ_m , K_{reg} , K_{ε} , and K_{ξ} such that the joint distribution $P^{X,\varepsilon}$ satisfies:

(i) $\lambda_{\min}(\mathbb{E}(XX')) \geq \lambda_m;$ (ii) $\mathbb{E}\left[||vec(XX' - \mathbb{E}[XX'])||^2\right] \leq K_{\text{reg}};$ (iii) $\mathbb{E}[||X\varepsilon||^4] \leq K_{\varepsilon};$ (iv) $\mathbb{E}[\xi^4] / \mathbb{E}[\xi^2]^2 \leq K_{\xi}.$ Assumption 5.2 defines a broad nonparametric class of distributions delineated by the different constants λ_m , K_{reg} , K_{ε} , and K_{ξ} . These constants appear explicitly in the construction of our CIs. In practice, the user needs to specify their values, which should be done with care. We elaborate on these choices in more details in Section 5.4. Relying on explicit constants may seem restrictive compared to asymptotic standard inference. However, as explained in the simpler case of Example 5.1, these constraints are somewhat inescapable to obtain nonasymptotic properties except if one is willing to adopt alternative parametric assumptions.

Overall, the different parts of Assumption 5.2 strengthen the moment conditions of the basic linear model Θ . Part (i) rules out $\mathbb{E}[XX']$ matrices arbitrarily close to being singular, an unfavorable situation in which β_0 is not identified. Although the theoretical parameter β_0 is defined on Θ , in general, the corresponding OLS estimator $\hat{\beta}$ is not properly defined with probability one in finite samples due to the possible singularity of $n^{-1}\sum_{i=1}^{n} X_i X'_i$. Part (ii) helps control the concentration of the latter empirical mean and ensures the proper definition of $\hat{\beta}$ with large probability for every (large enough) sample size. Parts (i) to (iii) are critical in the construction of our CI as they enable us to ensure $(u'\hat{V}u)^{-1/2}u'(\hat{\beta} - \beta_0)$ is "close" to the simpler-to-analyze quantity $\sum_{i=1}^{n} \xi_i / \sqrt{\sum_{i=1}^{n} \xi_i^2}$. More precisely, they first guarantee that the linearization

$$\sqrt{n}u'(\widehat{\beta} - \beta_0) \in \left[\frac{1}{\sqrt{n}}\sum_{i=1}^n \xi_i \pm R_{n,\text{lin}}(\delta)\right]$$

holds with probability at least $1 - \delta$, where

$$R_{n,\mathrm{lin}}(\delta) := \frac{\|u\|}{\lambda_m \lambda_{\mathrm{min}}(\overline{XX'})} \sqrt{\frac{2K_{\mathrm{reg}}}{n\delta}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4}$$

is an explicit linearization error term and $\overline{XX'}$ a short-hand notation for $\frac{1}{n}\sum_{i=1}^{n}X_iX'_i$.

Second, with probability at least $1-\delta$ too, they enable us to control the oracle estimator $\sum_{i=1}^{n} \xi_i^2$ of u'Vu by the feasible quantity $u'\hat{V}u + ||u||^2 R_{n,\text{var}}(\delta)$, where

$$\begin{aligned} R_{n,\text{var}}(\delta) &:= \frac{2}{n\lambda_m^2\lambda_{\min}(\overline{XX'})^2} \sqrt{\frac{K_{\varepsilon}}{\delta}} \times \frac{1}{n} \sum_{i=1}^n \|X_i\|^4 + \\ &\frac{2\sqrt{2}}{\lambda_m^2\lambda_{\min}(\overline{XX'})\sqrt{n}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4} \times \frac{1}{n} \sum_{i=1}^n \|X_i\|^3 |\widehat{\varepsilon}_i| + \\ &\frac{K_{\text{reg}}}{n\delta\lambda_m^2\lambda_{\min}(\overline{XX'})^2} \times \frac{1}{n} \sum_{i=1}^n \|X_i\|^2 |\widehat{\varepsilon}_i|^2 + \frac{2}{\lambda_m\lambda_{\min}(\overline{XX'})^2} \sqrt{\frac{K_{\text{reg}}}{n\delta}} \times \frac{1}{n} \sum_{i=1}^n \|X_i\|^2 |\widehat{\varepsilon}_i|^2. \end{aligned}$$

This bound is finite only when the latter matrix is invertible, a condition that we ensure when dealing with the proper definition of $\hat{\beta}$.

Parts (i) and (ii) of Assumption 5.2 have one major drawback: they are impacted by a translation of the covariate vector. As a result, the confidence intervals we build on (a combination of) the slope coefficients $(\beta_{0,1}, \ldots, \beta_{0,p})$ can be strongly affected by the location of the covariates. Fortunately, when one is interested in parameters $u'\beta_0$ that do not depend on the intercept of the linear model, i.e. the first entry in u is equal to zero, X can be replaced with $\widetilde{X} = (1, X_1 - \mathbb{E}[X_1], \ldots, X_p - \mathbb{E}[X_p])'$ in Assumption 5.2.

 $\sum_{i=1}^{n} \xi_i / \sqrt{\sum_{i=1}^{n} \xi_i^2}$ is a self-normalized sum (SNS) of *i.i.d.* variables. There exist several approaches to control the concentration of a SNS that we leverage to build our CI. They

use a bound on the kurtosis of ξ , the influence function of $\sqrt{n}u'(\hat{\beta} - \beta_0)$, which is set in Part (iv).

In fact, our confidence interval comprises two distinct regimes depending on the results used to control that SNS: an exponential bound for small sample sizes (Exp regime); a bound that involves Edgeworth expansions for larger n (Edg regime). In the latter regime, the bound can be improved under the next assumption, which deals with situations where we can assume some continuity with respect to Lebesgue's measure of the regressors or residuals, hence on the variable ξ .

Assumption 5.3 (Regularity assumption on the distribution of ξ). The parameter $\theta = (\beta_0, \mathbf{P}^{X,\varepsilon})$ belongs to Θ and $\mathbf{P}^{X,\varepsilon}$ is such that ξ has an absolutely continuous component with respect to Lebesgue's measure.

Except for quite peculiar combinations, the only practical case ruled out by this assumption is a situation with discrete covariates and discrete residuals at the same time; for instance, a linear probability model (support(Y) = {0,1}) without any continuous regressors. In particular, a residual ε with a continuous distribution with respect to Lebesgue's measure is enough to satisfy the condition. Under this condition, we derive tighter CIs with maintained nonasymptotic properties.

In the Edgeworth regime, we introduce the generic notation $\Delta_{n,E}$ that can correspond to two different quantities according to the maintained assumptions. In the "general" case, namely under Assumption 5.2 only, $\Delta_{n,E}$ stands for

$$\Delta_{n,E}^{\text{gen}} := \frac{0.2(K_{\xi}^{3/4} + 1)}{\sqrt{n}} + \frac{0.2K_{\xi} + 0.12K_{\xi}^{3/2}}{n} + r_n^{\text{gen}}$$

where $r_n^{\text{gen}} = O(n^{-5/4})$ is an explicit term given in Section 5.D.1 which depends on K_{ξ} and n only. In the "continuous" case, that is under Assumptions 5.2 and 5.3, $\Delta_{n,\text{E}}$ is equal to

$$\Delta_{n,E}^{\text{cont}} := \frac{0.195 \, K_{\xi} + 0.0147 K_{\xi}^{3/2}}{n} + r_n^{\text{cont}}$$

where $r_n^{\text{cont}} = O(n^{-5/4})$ is an explicit term given in Section 5.D.2. The remainder r_n^{cont} depends on K_{ξ} , n, and the tail behavior of the characteristic function of ξ through the quantity

$$\kappa := \sup_{t: |t| \ge 2t_1^* \pi / (K_{\xi}^{3/4} + 1)} |f(t)|$$

with f the characteristic function of the normalized variable $\xi/\sqrt{\mathbb{E}[\xi^2]}$, $t_1^* := \theta_1^*/(2\pi)$ and θ_1^* the unique root in $(0, 2\pi)$ of the equation $\theta^2 + 2\theta \sin(\theta) + 6(\cos(\theta) - 1) = 0$. As mentioned by Shevtsova (2010), $t_1^* \approx 0.64$. Remark that κ depends implicitly on the choice of K_{ξ} that has been made in Assumption 5.2.

Construction of the confidence interval

Henceforth we fix a desired nominal level $1 - \alpha \in (0, 1)$ and a sample size n. Note that the following expressions depend on two free parameters $\omega \in (0, 1)$ and a > 1. We do not indicate this dependence to lighten notations. We first define two intervals, centered at

 $u'\hat{\beta}$, that we combine later to form the final confidence interval

$$CI_{u}^{Edg}(1-\alpha,n) := \left[u'\widehat{\beta} \pm \frac{Q_{n}^{Edg}}{\sqrt{n}} \sqrt{u'\widehat{V}u + \|u\|^{2}R_{n,var}(\omega\alpha/2)} \right],$$
$$CI_{u}^{Exp}(1-\alpha,n) := \left[u'\widehat{\beta} \pm \frac{Q_{n}^{Exp}}{\sqrt{n}} \sqrt{u'\widehat{V}u + \|u\|^{2}R_{n,var}(\omega\alpha/2)} \right].$$

Those two intervals are similar to $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha, n)$ in Equation (5.5) with the addition of $||u||^2 R_{n,\operatorname{var}}(\omega\alpha/2)$ in the variance term and the "modified Gaussian quantiles"

$$Q_n^{\text{Edg}} := \sqrt{a} \, q_{\mathcal{N}(0,1)} \left(1 - \alpha/2 + \nu_n^{\text{Edg}} \right) + \nu_n^{\text{Approx}},$$
$$Q_n^{\text{Exp}} := \sqrt{2(1+a) \left(1 - \ln\left(\alpha/2 - \nu_n^{\text{Exp}}\right) \right)} + \nu_n^{\text{Approx}},$$

which depend on some perturbation terms $\nu_n^{\rm Edg}, \nu_n^{\rm Exp}, \nu_n^{\rm Approx}$ defined as

$$\nu_n^{\text{Edg}} := \nu_n^{\text{Exp}} + \Delta_{n,\text{E}},$$
$$\nu_n^{\text{Exp}} := \frac{\omega\alpha + \exp\left(-n(1-1/a)^2/(2K_{\xi})\right)}{2}$$
$$\nu_n^{\text{Approx}} := \frac{R_{n,\text{lin}}(\omega\alpha/2)}{\sqrt{u'\widehat{V}u + \|u\|^2 R_{n,\text{var}}(\omega\alpha/2)}}.$$

Note that ν_n^{Approx} is a random quantity since $R_{n,\text{lin}}$, $R_{n,\text{var}}$ and \hat{V} depend on the sample. Therefore, unlike $\text{CI}_u^{\text{Asymp}}(1-\alpha,n)$, the "modified Gaussian quantiles" are random quantities, which introduces another source of randomness into the length of both $\text{CI}_u^{\text{Edg}}(1-\alpha,n)$ and $\text{CI}_u^{\text{Exp}}(1-\alpha,n)$.

Remark that for the proper definition of Q_n^{Edg} (resp. Q_n^{Exp}), ν_n^{Edg} (resp. ν_n^{Exp}) needs to be strictly smaller than $\alpha/2$. Since $\Delta_{n,\text{E}}$ and ν_n^{Exp} are decreasing in n, this will arise for sufficiently large samples. Furthermore, given the expression of ν_n^{Edg} , the Exp regime becomes available before the Edg one. Finally, when $\nu_n^{\text{Exp}} \geq \alpha/2$, neither $\text{CI}_u^{\text{Edg}}(1-\alpha,n)$ nor $\text{CI}_u^{\text{Exp}}(1-\alpha,n)$ are feasible. In addition, the proper definitions of $\text{CI}_u^{\text{Edg}}(1-\alpha,n)$ and $\text{CI}_u^{\text{Exp}}(1-\alpha,n)$ require the technical condition $\omega > 2K_{\text{reg}}/(n\alpha\lambda_m^2)$. We impose the latter to build a large probability event on which $\overline{XX'}$ is invertible and $\hat{\beta}$ is well-defined. Therefore, if the latter condition is not satisfied or if $\nu_n^{\text{Exp}} \geq \alpha/2$, we adopt the convention of defining our final CI as the entire real line. Otherwise, the Exp regime applies as long as $\nu_n^{\text{Edg}} \geq \alpha/2$. Finally, the Edg regime becomes available and is selected when $\nu_n^{\text{Edg}} < \alpha/2$. To sum up

$$\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n) := \begin{cases} \mathbb{R} & \text{if } n \leq 2K_{\operatorname{reg}}/(\omega\alpha\lambda_{m}^{2}) \text{ or } \nu_{n}^{\operatorname{Exp}} \geq \alpha/2 \\ \operatorname{CI}_{u}^{\operatorname{Exp}}(1-\alpha,n) & \text{else if } \nu_{n}^{\operatorname{Edg}} \geq \alpha/2, \\ \operatorname{CI}_{u}^{\operatorname{Edg}}(1-\alpha,n) & \text{else.} \end{cases}$$

The first regime in the definition of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ provides little information on $u'\beta_{0}$. It prevails in particular when the sample size is too small to ensure that $\widehat{\beta}$ is well-defined with large probability. Another interpretation is that the sample is not large enough to build a CI with both controlled length and guaranteed coverage.

We are now in a position to study the theoretical properties of our confidence interval $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$.

5.3.3 Theoretical results

We present two theorems. In Theorem 5.1, we show $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ is nonasymptotically conservative for every data distribution indexed by $\theta \in \Theta$ which also satisfies Assumption 5.2 or Assumptions 5.2 and 5.3. In Theorem 5.2, we derive the asymptotic pointwise exactness of our CI. In the rest of the paper, we use \mathbb{P} instead of $\mathbb{P}_{\operatorname{P}^{\otimes n}}$ to lighten notation.

Theorem 5.1 (Nonasymptotic conservativeness). Under Assumption 5.2 ("general" case, $\Delta_{n,E} = \Delta_{n,E}^{\text{gen}}$) or under Assumptions 5.2 and 5.3 ("continuous" case, $\Delta_{n,E} = \Delta_{n,E}^{\text{cont}}$), for every n > 1, $\alpha \in (0, 1)$, a > 1, $\omega \in (0, 1)$, we have

$$\mathbb{P}\left(u'\beta_0 \in CI_u^{Fin}(1-\alpha,n)\right) \ge 1-\alpha.$$
(5.6)

Theorem 5.1 achieves one of our main goals, namely proving it is possible to build a CI which is nonasymptotically conservative for a large class of DGPs without imposing independence between X and ε or parametric restrictions on ε . We remind the reader that the above-constructed CI is truly informative ($\neq \mathbb{R}$) under the conditions $n > 2K_{\text{reg}}/(\omega\alpha\lambda_m^2)$ and $\nu_n^{\text{Exp}} < \alpha/2$. Unsurprisingly, we can see that the smaller α gets, the larger n has to be to satisfy the former constraints. Before turning to asymptotic exactness, we also remark that, as a corollary of Theorem 5.1, $\text{CI}_u^{\text{Fin}}(1-\alpha,n)$ is asymptotically conservative uniformly over the set of parameters θ satisfying Assumption 5.2 (plus Assumption 5.3 for the "continuous" case) for any fixed choice of (a, ω) .

To prove pointwise asymptotic exactness of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$, we have to consider sequences $(\omega_n)_n$ and $(a_n)_n$ for the free parameters and place specific constraints on these sequences. We need an additional moment condition compared to Assumption 5.2 as well. The theorem writes as follows:

Theorem 5.2 (Asymptotic pointwise exactness). Let $b_n := a_n - 1$ and assume that $b_n = o(1), b_n \sqrt{n} \to +\infty, \omega_n = o(1), n^{1/2} \omega_n^{3/4} \to +\infty, \mathbb{E}[||X||^4] < +\infty$, and that the assumptions of Theorem 5.1 hold. Then, for every $\alpha \in (0, 1)$, we have

$$\lim_{n \to \infty} \mathbb{P}\Big(u'\beta_0 \in CI_u^{Fin}(1-\alpha, n)\Big) = 1-\alpha.$$
(5.7)

Inspection of the proof in fact shows that for every distribution satisfying the conditions of the theorem

$$CI_{u}^{Fin}(1-\alpha,n) = \left[u'\hat{\beta} \pm \frac{q_{\mathcal{N}(0,1)}(1-\alpha/2) + o_{P}(1)}{\sqrt{n}} \sqrt{u'\hat{V}u + o_{P}(1)} \right].$$

This result actually extends to any distribution indexed by $\theta \in \Theta^{\text{Asymp}}$. Consequently, $\operatorname{CI}_{u}^{\text{Fin}}(1-\alpha, n)$ is nonasymptotically conservative over a subset of Θ , which is not granted for $\operatorname{CI}_{u}^{\text{Asymp}}(1-\alpha, n)$, and behaves (in a pointwise asymptotic sense) similarly to $\operatorname{CI}_{u}^{\text{Asymp}}(1-\alpha, n)$ over Θ^{Asymp} .

5.4 Practical considerations

In Assumption 5.2, lower or upper bounds have to be imposed on four moments (or functions thereof) of the distribution $P_{X,\varepsilon}$. While *a priori* choices for those bounds are natural in specific cases (see Remark 5.1 below), it remains often difficult to have intuition about the values of these bounds or how to choose them. Therefore, we propose to replace these bounds by estimates of the corresponding moments:

- (i) We replace λ_m with $\lambda_{\min}(n^{-1}\sum_{i=1}^n X_i X_i')$.
- (ii) Instead of K_{reg} we use

$$\frac{1}{n}\sum_{i=1}^{n} \left\| vec\left(X_iX_i' - \left(n^{-1}\sum_{j=1}^{n}X_jX_j'\right)\right) \right\|^2$$

(iii) $n^{-1} \sum_{i=1}^{n} ||X_i \hat{\varepsilon}_i||^4$ replaces K_{ε} .

(iv) $n \sum_{i=1}^{n} \widehat{\xi}_{i}^{4} / (\sum_{i=1}^{n} \widehat{\xi}_{i}^{2})^{2}$ is used in lieu of K_{ξ} , with $\widehat{\xi}_{i} := u' \left(n^{-1} \sum_{j=1}^{n} X_{j} X_{j}' \right)^{-1} X_{i} \widehat{\varepsilon}_{i}$.

Using plug-ins rather than deterministic bounds has one major drawback from a theoretical point of view: $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ is no longer formally valid in finite samples. However, by construction of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$, the latter is almost surely at least as large as $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha,n)$ for any sample size even when deterministic bounds are replaced with plug-ins. Furthermore, we show in several simulation experiments conducted in Section 5.5 that the plug-in version of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ is substantially larger than $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha,n)$ in moderate samples. In a sense, the use of plug-ins in our approach can be compared to the approximation error faced in practice when one uses an inference procedure based on simulations such as that of Chernozhukov et al. (2009) or DiCiccio and Romano (2017). To further control the impact of sampling uncertainty when using a plug-in version of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$, a practical possibility would be to multiply the latter by $(1 + M/\sqrt{n})$ for some positive user-chosen M.

Remark 5.1 (On the possibility to overcome plug-ins). Choosing reasonable values for K_{ξ} and K_{reg} without resorting to a plug-in strategy turns out to be possible. A large class of univariate distributions exhibits a bound of at most 9 on the kurtosis: Normal, Laplace, asymmetric Laplace, Logistic, Uniform, Student with at least five degrees of freedom, two-point symmetric mixture of Normals, Gumbel, hyperbolic secant, and skewed Normal. This class includes both symmetric and asymmetric distributions, some of which only have a few number of finite moments (Student distributions with few degrees of freedom). We investigate the impact of the choice $K_{\xi} = 9$ as an alternative to the plug-in approach in the simulations presented in Section 5.5. At the same time, K_{reg} can be chosen naturally when X is bounded. However the obtained value for K_{reg} is likely to be crude. We thus do not assess this possibility on simulations.

Finally, we insist that whenever the first entry of u is equal to zero, i.e. $u'\beta_0$ does not depend on the intercept of the linear model, our method can be improved in practice. In this situation, the remark made in Section 5.3.2 applies: X can be replaced with $\widetilde{X} := (1, X_1 - \mathbb{E}[X_1], \dots, X_p - \mathbb{E}[X_p])'$ in Assumption 5.2 and the plug-in versions of the bounds have to be modified accordingly. The appealing feature of these modified plug-in bounds is their invariance to a translation of the covariate vector. These observations apply in particular when one is interested in building confidence intervals for individual slope coefficients of the linear model. In our simulations, we focus on conducting inference on slope coefficients so that we make use of the aforementioned modification of the plug-in bounds.

5.5 Simulations

The simulations aim to assess the performance of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ under different scenarios and compare it to that of the standard $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha,n)$.

We consider the following model, without endogeneity,

$$Y = 2 + 1X_1 - 3X_2 + \varepsilon, \quad (X_1, X_2)' \sim \mathcal{N}(\mu, \Sigma),$$

with $\mu := (0.5, 5)', \quad \Sigma := \begin{pmatrix} 1 & \sqrt{2} \times 0.5 \\ \sqrt{2} \times 0.5 & 2 \end{pmatrix}.$

We focus on individual coefficients of a linear model. The results are very for the coefficients of X_1 and X_2 ; we focus here on the former $\beta_{0,2} = 1$, that is, we choose u = (0, 1, 0)'.

We consider four different scenarios for the conditional distributions of residuals:⁹

$$- P_{\epsilon|X} = \mathcal{N} ormal(0, |X_1 + X_2|); - P_{\epsilon|X} = \mathcal{L}aplace(0, \sqrt{0.5|X_1 + X_2|}; - P_{\epsilon|X} = \mathcal{S}tudent(3\mathbb{1}\{X_1 + X_2 < \mathbb{E}[X_1 + X_2]\} + 4\mathbb{1}\{X_1 + X_2 \ge \mathbb{E}[X_1 + X_2]\}); - P_{\epsilon|X} = \mathcal{G}umbel(-\sqrt{(6/\pi^2)|X_1 + X_2|}\gamma, \sqrt{(6/\pi^2)|X_1 + X_2|})$$
 where $\gamma \approx 0.58$ is the Euler-Mascheroni constant.

We look a priori at sample sizes n ranging from one hundred to 25,000 to explore the behavior in small, moderate, and large samples. The distributions chosen for the residuals respect mean-independence and allow for heteroskedasticity.

Remember that our CI $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ requires known values for the bounds λ_{m} , K_{ε} , K_{reg} , K_{ξ} , as well as choices for the free parameters ω and α . As explained in the previous section, we use a plug-in strategy for the first three bounds. Regarding K_{ξ} , a natural bound of 9 is available or a plug-in estimate can be used.

The free parameters are chosen so as to satisfy the conditions of Theorem 5.2: $\omega = \omega_n = n^{-1/8}$, $a = a_n = 1 + n^{-1/4}$.

We compute $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha, n)$ under four different assumptions: the "general case" without additional regularity assumption (Assumption 5.2 only) and the "continuous case" with additional regularity assumption (with additional Assumption 5.3).¹⁰ The other two assumptions come from the possibility to improve the bound $\Delta_{n,\mathrm{E}}$ when assuming an unskewed distribution for ξ , which is indeed the case in the first three data-generating processes we consider We refer to Derumigny et al. (2021) for further details. In the simulations, the improvement appears negligible, and we focus on comparing the general and the continuous case.

Figure 5.1 show the realizations of $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha, n)$ and $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha, n)$ for the Gaussian data-generating process using the bound $K_{\xi} = 9$ (averages over 100 simulations).¹¹ Figure 5.2 is similar using a plug-in for K_{ξ} . They call for several comments. We begin with the case $K_{\xi} = 9$.

⁹The choice made for the scale parameter of the Laplace enables to have the same variance $(X_1 + X_2)^2$ as in the Normal case; idem for the location and scale parameter of the Gumbel, set such that it is centered with variance $(X_1 + X_2)^2$.

¹⁰In the continuous case, we set $\kappa = 0.99$ to compute $\Delta_{n,E}^{\text{cont}}$ and $\operatorname{CI}_{u}^{\text{Fin}}(1-\alpha,n)$. Derumigny et al. (2021) discusses the choice of this parameter; 0.99 can be deemed a rather conservative, encompassing a wide family of possible distribution for ξ .

¹¹The other data-generating processes leads to similar qualitative conclusions and are reported in Appendix 5.E.



Figure 5.1: $K_{\xi} = 9 - \alpha = 0.05 - P^{\varepsilon \mid X}$ Gaussian.

Figure 5.2: K_{ξ} by plug-in – $\alpha = 0.05 - P^{\varepsilon | X}$ Gaussian.



First, it happens that the condition to avoid the \mathbb{R} regime of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ requires quite large sample sizes to hold in practice. In fact, the graphs starts for $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ at $n \approx 5,500$ because, before $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n) = \mathbb{R}$. This result is disappointing for the practical use of our CI. Closer inspection reveals the plug-in estimates of K_{reg} (respectively λ_m) can be quite large (respectively close to 0), hence the necessity to have large n to satisfy $n > 2K_{\operatorname{reg}}/(\omega\alpha\lambda_m^2)$ (the first condition needed to exit the \mathbb{R} regime). A possible improvement to soften this constraint would be to rely on normalized moments, which may have a smaller plug-in counterpart in practice. We leave this possibility for future work.

The choice of ω also matters in this condition: the lower ω , the more difficult the condition is to hold; hence the more difficult it is to have an informative $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$. On the other hand, lower ω increases precision by reducing $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$'s length. Yet, given the simulation results, the priority is avoiding the \mathbb{R} regime. This explains the choice

of $\omega = n^{-1/8}$: it satisfies the conditions for $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha, n)$ to be pointwise asymptotically exact, but the rate towards 0 is slow. Although less influential, the choice of *a* displays a similar trade-off through the second condition that needs to be satisfied to avoid the \mathbb{R} regime: $\nu_{n}^{\operatorname{Exp}} < \alpha/2$.

For n large enough, $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ becomes informative in the sense that it enters the Exp, then Edg regime. However, the simulation reveals it remains far broader than $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha,n)$, and its coverage is equal to 1. As expected, the additional regularity assumption improves the precision of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$. They also enable a switch from the Exp to the Edg regime for smaller sample sizes. This shift is well visible in Figure 5.1 in the "general case" for $n \approx 10,000$. Finally, the simulations illustrate the pointwise asymptotic exactness of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$: its length tends to the one of $\operatorname{CI}_{u}^{\operatorname{Asymp}}(1-\alpha,n)$ when the sample size goes to infinity.

In Figure 5.2, K_{ξ} is replaced by a plug-in estimate. The comparison with the previous case shows this implies an even larger sample size to exit the \mathbb{R} regime in the continuous case assumption and, to a lesser extent, in the general case (in other DGP reported in the appendix, this effect of using a plug-in instead of the bound 9 on K_{ξ} is more visible).

5.6 Extension: linear regression with endogeneity

5.6.1 Model and standard Anderson-Rubin inference

We extend our analysis to linear regressions with endogeneity. We focus on the case of one scalar endogenous regressor D, one excluded instrument Z, and additional exogenous covariates G, which is a frequent case in practice. Throughout this section, we assume we observe $n \ i.i.d.$ observations $(D_i, G_i, Z_i, Y_i)_{i=1}^n$ that satisfy the following linear instrumental model.

Assumption 5.4. $(D_i, G_i, Z_i, Y_i)_{i=1}^n \stackrel{i.i.d.}{\sim} \mathbb{P}^{D,G,Z,Y}$ where $X = (1, D, G')' \in \mathbb{R}^p$, $D \in \mathbb{R}$, and $W = (1, Z, G')' \in \mathbb{R}^p$, $Z \in \mathbb{R}$, such that, for some random variable ε and vector β_0 :

$$Y = X'\beta_0 + \varepsilon = \beta_{0,1} + \beta_{0,2}D + G'\beta_{0,3:p} + \varepsilon,$$

and $\mathbb{E}[W\varepsilon] = 0, \mathbb{E}[||W||^2] < +\infty, \mathbb{E}[||W||^2\varepsilon^2] < +\infty, \lambda_{\min}(\mathbb{E}[WW']) > 0.$

The parameter set of the associated statistical model is

$$\Theta_{IV} := \{ \theta = (\beta_0, \mathcal{P}_{D,G,Z,\varepsilon}) \in \mathbb{R}^p \times \mathcal{P}_{D,G,Z,\varepsilon} \}$$

with $\mathcal{P}_{D,G,Z,\varepsilon}$ corresponding to

$$\left\{ \mathcal{P}_{D,G,Z,\varepsilon} \in \mathcal{P}(\mathbb{R}^{p+3}) : \mathbb{E}[W\varepsilon] = 0, \ \mathbb{E}[\|W\|^2] < +\infty, \ \mathbb{E}[\|W\|^2\varepsilon^2] < +\infty, \ \lambda_{\min}(\mathbb{E}[WW']) > 0 \right\}.$$

In such settings, the main parameter of interest is often the coefficient $\beta_{0,2}$ of the endogenous regressor. We thus focus on constructing confidence sets for $\beta_{0,2}$. In the definition of this model, we do not assume that $\mathbb{E}[ZW']$ has rank p, allowing for challenging situations in which β_0 is not identified.

Our CS addresses two potential issues. First, similarly to the OLS case, asymptotic approximation might be questioned and we construct a CS with finite-sample guaranteed coverage. Second, we want to conduct inference on $\beta_{0,2}$ even when the latter coefficient

might not be identified, a situation that arises when the instrument is weak, i.e. when the instrument is not correlated with the endogenous variable.

In order to do so, we adapt the well-known Anderson and Rubin (1949) construction. For every $\tilde{\beta}_2 \in \mathbb{R}$, we consider the random variable $Y - \tilde{\beta}_2 D$. We orthogonally project this random variable on the space generated by the linear combinations of W, giving the decomposition

$$Y - \tilde{\beta}_2 D = \alpha(\tilde{\beta}_2) + \lambda(\tilde{\beta}_2) Z + G' \mu(\tilde{\beta}_2) + \zeta(\tilde{\beta}_2),$$

where $\alpha(\tilde{\beta}_2)$, $\lambda(\tilde{\beta}_2)$, and $\mu(\tilde{\beta}_2)$ denote the coefficients of the projection for the different components of W and $\zeta(\tilde{\beta}_2)$ is a random error term with expectation 0 and uncorrelated with W. Note that we choose to make explicit the dependence on the parameter $\tilde{\beta}_2$ in the previous decomposition. Moreover, using the assumption that $\mathbb{E}[WW']^{-1}$ exists, the function $\psi(\tilde{\beta}_2) := (\alpha(\tilde{\beta}_2), \lambda(\tilde{\beta}_2), \mu(\tilde{\beta}_2)')'$ has the explicit form $\mathbb{E}[WW']^{-1}\mathbb{E}[W(Y - \tilde{\beta}_2 D)]$.

For the choice $\tilde{\beta}_2 = \beta_{0,2}$, we can write $Y - D\tilde{\beta}_2 = W'\psi_0$ with $\psi_0 := (\beta_{0,1}, 0, \beta'_{0,3})'$, implying that $\psi(\beta_{0,2}) = \psi_0$. For every $\tilde{\beta}_2 \in \mathbb{R}$, let $T_n(\tilde{\beta}_2)$ denote the Wald-statistic for testing the hypothesis $H_0: \lambda(\tilde{\beta}_2) = 0$ against $H_1: \lambda(\tilde{\beta}_2) \neq 0$,

$$T_n(\tilde{\beta}_2) = \frac{\sqrt{n} \, \mathbf{e}_2' \hat{\psi}(\tilde{\beta}_2)}{\sqrt{\mathbf{e}_2' \hat{V}(\tilde{\beta}_2) \mathbf{e}_2}},$$

with $\widehat{\psi}(\widetilde{\beta}_2)$ the OLS estimator of $\psi(\widetilde{\beta}_2)$ given by

$$\widehat{\psi}(\widetilde{\beta}_2) := \left(\sum_{i=1}^n W_i W_i'\right)^{-1} \left(\sum_{i=1}^n W_i (Y_i - \widetilde{\beta}_2 D_i)\right),$$

the estimator $\hat{V}(\tilde{\beta}_2)$ of the asymptotic variance of $\hat{\psi}(\tilde{\beta}_2)$ given by

$$\widehat{V}(\widetilde{\beta}_2) := \left(\frac{1}{n} \sum_{i=1}^n W_i W_i'\right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n W_i W_i' \widehat{\zeta}_i (\widetilde{\beta}_2)^2\right) \left(\frac{1}{n} \sum_{i=1}^n W_i W_i'\right)^{-1},$$

and the residuals $\hat{\zeta}_i(\tilde{\beta}_2) := Y_i - \tilde{\beta}_2 D_i - W'_i \hat{\psi}(\tilde{\beta}_2)$. Remark that $\mathbf{e}'_2 \hat{\psi}(\tilde{\beta}_2)$ can be recognized as the OLS estimator of $\lambda(\tilde{\beta}_2)$ and $\mathbf{e}'_2 \hat{V}(\tilde{\beta}_2)\mathbf{e}_2$ as its estimated asymptotic variance. Note that, since $\mathbb{E}[WW']$ is supposed to be invertible, $\hat{\psi}(\tilde{\beta}_2)$ and $\hat{V}(\tilde{\beta}_2)$ are well-defined with large probability for *n* large enough (see Section 5.3 for more details).

When $\tilde{\beta}_2 = \beta_{0,2}, T_n(\tilde{\beta}_2) \xrightarrow{d} \mathcal{N}(0,1)$. This yields the classical Anderson-Rubin (AR) confidence set for $\beta_{0,2}$ defined by

$$\operatorname{CS}^{AR,as}(1-\alpha,n) := \left\{ \widetilde{\beta}_2 \in \mathbb{R} : |T_n(\widetilde{\beta}_2)| \le q_{\mathcal{N}(0,1)}(1-\alpha/2) \right\},$$
(5.8)

which is asymptotically exact. This CS has been shown to be asymptotically uniformly exact over large classes of data-generating process and to remain valid even when the instrument is weak (see Andrews et al. (2019) for a recent survey of the results in the weak IV literature). Moreover, by adjusting for the number of degrees of freedom in the regression, this procedure is known to be of exact level in finite samples under the assumption that $P_{\varepsilon,\nu|Z} \sim \mathcal{N}(\mathbf{0}_2, \Sigma)$, with Σ known (Moreira (2003)) and with ν the error term in a first stage regression of D on W. To the best of our knowledge, we are the first to propose a nonasymptotic confidence set in a linear IV model without imposing that the residuals are Normal or independent from W.

5.6.2 Our confidence set

We introduce the random variable

$$\xi := \mathbf{e}_2' \mathbb{E}[WW']^{-1} W \varepsilon$$

that comes from the linearization of $\hat{\psi}(\hat{\beta}_2)$ used in the construction of our nonasymptotic CS. We present direct analogues of Assumptions 5.2 and 5.3 which help us derive the nonasymptotic validity of our CS on a subset of Θ_{IV} . Intuition behind the "IV-version" of Assumptions 5.2 and 5.3 is immediate given the detailed description of the assumptions presented in Section 5.3.

Assumption 5.5 (Bounds on DGP). The parameter $\theta = (\beta_0, P_{D,G,Z,\varepsilon})$ belongs to Θ_{IV} and there exist positive constants λ_m , K_{reg} , K_{ε} , and K_{ξ} such that the joint distribution $P^{X,\varepsilon}$ satisfies:

(i) $\lambda_{\min}(\mathbb{E}(WW')) \geq \lambda_m;$ (ii) $\mathbb{E}\left[||vec(WW' - \mathbb{E}[WW'])||^2\right] \leq K_{reg};$ (iii) $\mathbb{E}[||W\varepsilon||^4] \leq K_{\varepsilon};$ (iv) $\mathbb{E}\left[\xi^4\right] / \mathbb{E}\left[\xi^2\right]^2 \leq K_{\xi}.$

Assumption 5.6 (Regularity assumption on the distribution of ξ). The parameter $\theta = (\beta_0, P^{D,G,Z,\varepsilon})$ belongs to Θ_{IV} and $P^{W,\varepsilon}$ is such that ξ has an absolutely continuous component with respect to Lebesgue's measure.

We can now define two (random) sets that we combine to form the final confidence set

$$CS^{Edg}(1-\alpha,n) := \left\{ \widetilde{\beta}_2 : |\mathbf{e}_2'\widehat{\psi}(\widetilde{\beta}_2)| \le \frac{Q_n^{Edg}(\widetilde{\beta}_2)}{\sqrt{n}} \sqrt{\mathbf{e}_2'\widehat{V}(\widetilde{\beta}_2)\mathbf{e}_2 + R_{n,var}(\omega\alpha/2,\widetilde{\beta}_2)} \right\},\\ CS^{Exp}(1-\alpha,n) := \left\{ \widetilde{\beta}_2 : |\mathbf{e}_2'\widehat{\psi}(\widetilde{\beta}_2)| \le \frac{Q_n^{Exp}(\widetilde{\beta}_2)}{\sqrt{n}} \sqrt{\mathbf{e}_2'\widehat{V}(\widetilde{\beta}_2)\mathbf{e}_2 + R_{n,var}(\omega\alpha/2,\widetilde{\beta}_2)} \right\}.$$

Those two sets are very close to $CS^{AR,as}(1-\alpha,n)$ in Equation (5.8). There are two main differences: estimated variance is inflated by $R_{n,var}(\omega\alpha/2,\tilde{\beta}_2)$ and the Gaussian quantile $q_{\mathcal{N}(0,1)}(1-\alpha/2)$ is replaced with either of the following "modified Gaussian quantiles"

$$Q_n^{\text{Edg}}(\widetilde{\beta}_2) := \sqrt{a} \, q_{\mathcal{N}(0,1)} \Big(1 - \alpha/2 + \nu_n^{\text{Edg}} \Big) + \nu_n^{\text{Approx}}(\widetilde{\beta}_2),$$
$$Q_n^{\text{Exp}}(\widetilde{\beta}_2) := \sqrt{2(1+a) \Big(1 - \ln\left(\alpha/2 - \nu_n^{\text{Exp}}\right) \Big)} + \nu_n^{\text{Approx}}(\widetilde{\beta}_2),$$

which depend on some perturbation terms $\nu_n^{\text{Edg}}, \nu_n^{\text{Exp}}, \nu_n^{\text{Approx}}(\tilde{\beta}_2)$ defined as

$$\nu_n^{\text{Edg}} := \nu_n^{\text{Exp}} + \Delta_{n,\text{E}},$$
$$\nu_n^{\text{Exp}} := \frac{\omega\alpha + \exp\left(-n(1-1/a)^2/(2K_{\xi})\right)}{2},$$
$$\nu_n^{\text{Approx}}(\tilde{\beta}_2) := \frac{R_{n,\text{lin}}(\omega\alpha/2)}{\sqrt{\mathbf{e}_2'\hat{V}(\tilde{\beta}_2)\mathbf{e}_2 + R_{n,\text{var}}(\omega\alpha/2,\tilde{\beta}_2)}}.$$

The terms $R_{n,\text{lin}}(\delta)$ and $R_{n,\text{var}}(\delta, \tilde{\beta}_2)$ are again direct extensions of the corresponding terms in the exogenous case:

$$R_{n,\mathrm{lin}}(\delta) := \frac{1}{\lambda_m \lambda_{\mathrm{min}}(\overline{WW'})} \sqrt{\frac{2K_{\mathrm{reg}}}{n\delta}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4}$$

and

$$\begin{aligned} R_{n,\text{var}}(\delta,\tilde{\beta}_{2}) &:= \frac{2}{n\lambda_{m}^{2}\lambda_{\min}(\overline{WW'})^{2}}\sqrt{\frac{K_{\varepsilon}}{\delta}} \times \frac{1}{n}\sum_{i=1}^{n}\|W_{i}\|^{4} \\ &+ \frac{2\sqrt{2}}{\lambda_{m}^{2}\lambda_{\min}(\overline{WW'})\sqrt{n}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4} \times \frac{1}{n}\sum_{i=1}^{n}\|W_{i}\|^{3}|\hat{\zeta}_{i}(\tilde{\beta}_{2})|^{2} \\ &+ \frac{K_{\text{reg}}}{n\delta\lambda_{m}^{2}\lambda_{\min}(\overline{WW'})^{2}} \times \frac{1}{n}\sum_{i=1}^{n}\|W_{i}\hat{\zeta}_{i}(\tilde{\beta}_{2})\|^{2} \\ &+ \frac{2}{\lambda_{m}\lambda_{\min}(\overline{WW'})^{2}}\sqrt{\frac{K_{\text{reg}}}{n\delta}} \times \frac{1}{n}\sum_{i=1}^{n}\|W_{i}\hat{\zeta}_{i}(\tilde{\beta}_{2})\|^{2}. \end{aligned}$$

In line with the exogenous case, we have to set our CS equal to the entire real line when n is such that $2K_{\rm reg}/(\omega\alpha\lambda_m^2)$ or $\nu_n^{\rm Exp} \ge \alpha/2$. For "moderate" (resp. "large") sample sizes, we use $\rm CS^{\rm Exp}(1-\alpha, n)$ (resp. $\rm CS^{\rm Edg}(1-\alpha, n)$) as our confidence set. To sum up

$$\mathrm{CS}^{\mathrm{Fin}}(1-\alpha,n) := \begin{cases} \mathbb{R} & \text{if } n \leq 2K_{\mathrm{reg}}/(\omega\alpha\lambda_m^2) \text{ or } \nu_n^{\mathrm{Exp}} \geq \alpha/2, \\ \mathrm{CS}^{\mathrm{Exp}}(1-\alpha,n) & \text{else if } \nu_n^{\mathrm{Edg}} \geq \alpha/2, \\ \mathrm{CS}^{\mathrm{Edg}}(1-\alpha,n) & \text{else.} \end{cases}$$

5.6.3 Theoretical results

Under the assumptions introduced in the previous subsection, we are able to prove the nonasymptotic conservativeness as well as the asymptotic pointwise exactness of $CS^{Fin}(1-\alpha, n)$.

Theorem 5.3 (Nonasymptotic conservativeness). Under Assumption 5.5 ("general" case, $\Delta_{n,E} = \Delta_{n,E}^{\text{gen}}$) or under Assumptions 5.5 and 5.6 ("continuous" case, $\Delta_{n,E} = \Delta_{n,E}^{\text{cont}}$), for every n > 1, $\alpha \in (0, 1)$, a > 1, $\omega \in (0, 1)$, we have

$$\mathbb{P}\Big(\beta_{0,2} \in CS^{Fin}(1-\alpha,n)\Big) \ge 1-\alpha.$$
(5.9)

Theorem 5.4 (Asymptotic pointwise exactness). Let $b_n := a_n - 1$ and assume that $b_n = o(1), b_n \sqrt{n} \to +\infty, \omega_n = o(1), n^{1/2} \omega_n^{3/4} \to +\infty, \mathbb{E}[||W||^4] < +\infty$, and that the assumptions of Theorem 5.3 hold. Then, for every $\alpha \in (0, 1)$, we have

$$\lim_{n \to \infty} \mathbb{P}\Big(\beta_{0,2} \in CS^{Fin}(1-\alpha, n)\Big) = 1-\alpha.$$
(5.10)

Beyond pointwise asymptotic exactness, we can also ensure $CS^{Fin}(1 - \alpha, n)$ is asymptotically uniformly conservative over the set of parameters θ satisfying Assumption 5.5 (plus Assumption 5.6 for the "continuous" case) for any fixed choice of (a, ω) .

References

- Acemoglu, D., Ozdaglar, A. and Tahbaz-Salehi, A. (2017), 'Microeconomic origins of macroeconomic tail risks', *American Economic Review* 107(1), 54–108.
- Anderson, T. W. and Rubin, H. (1949), 'Estimation of the parameters of a single equation in a complete system of stochastic equations', *The Annals of Mathematical Statistics* 20(1), 46–63.
- Andrews, D. W. K. and Guggenberger, P. (2019), 'Identification- and singularity-robust inference for moment condition models', *Quantitative Economics* **10**(4), 1703–1746.
- Andrews, I. and Mikusheva, A. (2016), 'Conditional inference with a functional nuisance parameter', *Econometrica* 84(4), 1571–1612.
- Andrews, I., Stock, J. H. and Sun, L. (2019), 'Weak instruments in instrumental variables regression: Theory and practice', Annual Review of Economics 11, 727–753.
- Armstrong, T. B. and Kolesár, M. (2020), 'Simple and honest confidence intervals in nonparametric regression', *Quantitative Economics* 11(1), 1–39.
- Ascari, G., Fagiolo, G. and Roventini, A. (2015), 'Fat-tail distributions and business-cycle models', *Macroeconomic Dynamics* 19(2), 465–476.
- Bahadur, R. R. and Savage, L. J. (1956), 'The nonexistence of certain statistical procedures in nonparametric problems', *The Annals of Mathematical Statistics* **27**(4), 1115–1122.
- Bentkus, V. (2003), 'On the dependence of the berry–esseen bound on dimension', *Journal* of Statistical Planning and Inference **113**(2), 385 402.
- Bentkus, V. and Götze, F. (1996), 'The berry-esseen bound for student's statistic', Ann. Probab. 24(1), 491–503.
- Bentkus, V. K., Bloznelis, M. and Götze, F. (1996), 'A berry-esseen bound for student's statistic in the non-i.i.d case', *Journal of Theoretical Probability* 9, 765–796.
- Bentkus, V. K. and Dzindzalieta, D. (2015), 'A tight gaussian bound for weighted sums of rademacher random variables', *Bernoulli* **21**(2), 1231–1237.
- Berry, A. (1941), 'The accuracy of the gaussian approximation to the sum of independent variates', *Transactions of the American Mathematical Society* **49**, 122–136.
- Bertail, P., Gautherat, E. and Harari-Kermadec, H. (2008), 'Exponential bounds for multivariate self-normalized sums', *Electron. Commun. Probab.* **13**, 628–640.
- Bertanha, M. and Moreira, M. J. (2016), 'Impossible inference in econometrics: Theory and applications'.
- Boucheron, S., Lugosi, G. and Massart, P. (2013), *Concentration inequalities: A nonasymp*totic theory of independence, Oxford university press.
- Chernozhukov, V., Hansen, C. and Jansson, M. (2009), 'Finite sample inference for quantile regression models', *Journal of Econometrics* **152**(2), 93 103. Nonparametric and Robust Methods in Econometrics.

- Derumigny, A., Girard, L. and Guyonvarch, Y. (2021), 'Explicit non-asymptotic bounds for the distance to the first-order edgeworth expansion', *arXiv preprint arXiv:2101.05780*.
- DiCiccio, C. J. and Romano, J. P. (2017), 'Robust permutation tests for correlation and regression coefficients', *Journal of the American Statistical Association* **112**(519), 1211–1220.
- Dufour, J.-M. (1997), 'Some impossibility theorems in econometrics with applications to structural and dynamic models', *Econometrica: Journal of the Econometric Society* pp. 1365–1387.
- Efron, B. (1969), 'Student's t-test under symmetry conditions', *Journal of the American Statistical Association* **64**(328), 1278–1302.
- Esseen, C.-G. (1942), 'On the liapunoff limit of error in the theory of probability', Arkiv för Matematik, Astronomi och Fysik. .
- Gossner, O. and Schlag, K. H. (2013), 'Finite-sample exact tests for linear regressions with bounded dependent variables', *Journal of Econometrics* **177**(1), 75–84.
- Hall, P. and Jing, B.-Y. (1995), 'Uniform coverage bounds for confidence intervals and berry-esseen theorems for edgeworth expansion', Ann. Statist. 23(2), 363–375.
- Hogben, L. (2006), Handbook of linear algebra, CRC press.
- Kasy, M. (2019), 'Uniformity and the delta method', *Journal of Econometric Methods* 8(1).
- Kleibergen, F. (2002), 'Pivotal statistics for testing structural parameters in instrumental variables regression', *Econometrica* **70**(5), 1781–1803.
- Kuchibhotla, A. K., Brown, L. D., Buja, A. and Cai, J. (2019), 'All of linear regression'.
- Kuchibhotla, A. K., Brown, L. D., Buja, A., George, E. I. and Zhao, L. (2020), 'Valid post-selection inference in model-free linear regression', *Annals of Statistics*.
- Moreira, M. J. (2003), 'A conditional likelihood ratio test for structural models', *Econometrica* **71**(4), 1027–1048.
- Pinelis, I. (1994), 'Extremal probabilistic problems and hotelling's t^2 test under a symmetry condition', Ann. Statist. **22**(1), 357–368.
- Pinelis, I. (2011a), 'On the berry-esseen bound for the student statistic'.
- Pinelis, I. (2011b), 'Relations between the first four moments'.
- Raič, M. (2019), 'A multivariate berry–esseen theorem with explicit constants', Bernoulli 25(4A), 2824–2853.
- Romano, J. P. and Wolf, M. (2000), 'Finite sample nonparametric inference and large sample efficiency', *The Annals of Statistics* **28**(3), 756–778.

- Schreuder, N., Brunel, V.-E. and Dalalyan, A. (2020), A nonasymptotic law of iterated logarithm for general m-estimators, in S. Chiappa and R. Calandra, eds, 'Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics', Vol. 108 of *Proceedings of Machine Learning Research*, PMLR, pp. 1331–1341.
- Shao, Q.-M. (2005), An explicit Berry-Esseen bound for Student's t-statistic via Stein's method, Lecture Notes Series, Institute for Mathematical Sciences, National University of Singapore, pp. 143–155.
- Shevtsova, I. (2010), 'Refinement of estimates for the rate of convergence in lyapunov's theorem', *Dokl. Akad. Nauk* **435**(1), 26–28.
- Student (1908), 'The probable error of a mean', Biometrika 6(1), 1-25.
- Tuvaandorj, P. (2020), Robust permutation tests in linear instrumental variables regression, Technical report.
- van der Vaart, A. (2000), Asymptotics Statistics, Cambridge University Press.

Appendix 5.A Proof of results in Section 5.3

5.A.1 Nonasymptotic conservativeness for the Edgeworth regime

Lemma 5.1. For every $\alpha \in (0,1)$, a > 1, $n \ge 1$ and $\omega \in (0,1)$, if $n\omega > 2K_{\text{reg}}/(\alpha \lambda_m^2)$ and $\nu_n^{\text{Edg}} < \alpha/2$, we have

$$\mathbb{P}\left(\left|u'\widehat{\beta}-u'\beta_{0}\right| \leq Q_{n}^{\mathrm{Edg}}n^{-1/2}\sqrt{u'\widehat{V}u+\left\|u\right\|^{2}R_{n,\mathrm{var}}(\omega\alpha/2)}\right) \geq 1-\alpha$$

Proof of Lemma 5.1. Let us define $\delta := \omega \alpha/2$. We want to show

$$\mathbb{P}\left(\frac{|u'\widehat{\beta} - u'\beta_0|}{\sqrt{u'\widehat{V}u + \|u\|^2 R_{n,\text{var}}(\delta)}} > Q_n^{\text{Edg}} n^{-1/2}\right) \le \alpha.$$
(5.11)

The proof is divided in three steps. In the first two ones, we derive two key intermediary inequalities that hold with high probability. In the final one, we combine those building bricks to obtain our result.

Step 1. Control of variance and linearization In this first step, we determine a high probability event (Step 1a.) on which we are able to control the residual term coming from the linearization of $u'\hat{\beta}$ (Step 1b.) and upper bound the oracle variance by a feasible quantity (Step 1c.).

Step 1a. Finding an event \mathcal{A} of high probability on which Steps 1b and 1c hold. Combining Assumption 5.2 (ii) and Lemma 5.5 with $A_i := X_i X'_i$, for $i = 1, \ldots, n$, we obtain

$$\mathbb{P}\left(\underbrace{\left\|\frac{1}{n}\sum_{i=1}^{n}X_{i}X_{i}'-\mathbb{E}[XX']\right\| \leq \sqrt{\frac{K_{\text{reg}}}{n\delta}}}_{=:\mathcal{A}_{1}}\right) \geq 1-\delta.$$
(5.12)

On the event \mathcal{A}_1 , thanks to Lemma 5.6, we have

$$\lambda_{\min}(\overline{XX'}) \ge \lambda_{\min}(\mathbb{E}(XX')) - \sqrt{\frac{K_{\text{reg}}}{n\delta}},$$
(5.13)

where $\overline{XX'}$ is a short-hand notation for $\frac{1}{n}\sum_{i=1}^{n}X_iX'_i$. Combining the constraint on ω , which ensures $n\omega \geq 2K_{\text{reg}}/(\alpha\lambda_m^2)$, with Assumption 5.2(i), we get on \mathcal{A}_1

$$\lambda_{\min}(\overline{XX'}) \ge \lambda_m - \sqrt{\frac{K_{\text{reg}}}{n\delta}} > 0, \qquad (5.14)$$

which implies that $\overline{XX'}$ is invertible.

In parallel, Assumption 5.2 (iii) and Lemma 5.7 provide the following inequality

$$\mathbb{P}\left(\underbrace{\left\|\frac{1}{n}\sum_{i=1}^{n}X_{i}\varepsilon_{i}\right\| \leq \sqrt{\frac{2}{n}\left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4}}}_{=:\mathcal{A}_{2}}\right) \geq 1-\delta.$$
(5.15)

Finally, we define $\mathcal{A} := \mathcal{A}_1 \cap \mathcal{A}_2$, which satisfies $\mathbb{P}(\mathcal{A}) \geq 1-2\delta$ thanks to Equations (5.12) and (5.15) and the law of total probability. Note that on \mathcal{A} , the estimator $\hat{\beta}$ is well-defined since the matrix $\overline{XX'}$ is invertible.

Step 1b. Linearization. The goal in this step is to formalize the approximation $u'\hat{\beta} \approx u'\beta_0 + \overline{\xi}_n$ where $\overline{\xi}_n$ has the following expression

$$\overline{\xi}_n := \frac{1}{n} \sum_{i=1}^n \xi_i = \frac{1}{n} \sum_{i=1}^n u' \mathbb{E}[XX']^{-1} X_i \varepsilon_i.$$

On the event \mathcal{A} , we will show that

$$\sqrt{n} \left| u'\hat{\beta} - u'\beta_0 - \overline{\xi}_n \right| \le \frac{\|u\|}{\lambda_m \lambda_{\min}(\overline{XX'})} \sqrt{\frac{2K_{\text{reg}}}{n\delta}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4} =: R_{n,\text{lin}}(\delta).$$
(5.16)

Cauchy-Schwarz's inequality yields

$$\begin{aligned} \left| u'\widehat{\beta} - u'\beta_0 - \overline{\xi}_n \right| &\leq \left\| u \right\| \left\| \widehat{\beta} - \beta_0 - \mathbb{E}[XX']^{-1} \frac{1}{n} \sum_{i=1}^n X_i \varepsilon_i \right\| \\ &= \left\| u \right\| \left\| \mathbb{E}[XX']^{-1} \left(\frac{1}{n} \sum_{i=1}^n X_i X'_i - \mathbb{E}[XX'] \right) (\widehat{\beta} - \beta_0) \right\|, \end{aligned}$$
(5.17)

where the second line comes from the following computation (using the equality $\hat{\beta} - \beta_0 = (n^{-1} \sum_{i=1}^n X_i X'_i)^{-1} n^{-1} \sum_{i=1}^n X_i \varepsilon_i)$

$$\begin{aligned} \left\| \mathbb{E}[XX']^{-1} \left(\frac{1}{n} \sum_{i=1}^{n} X_i X'_i - \mathbb{E}[XX'] \right) (\widehat{\beta} - \beta_0) \right\| \\ &= \left\| \mathbb{E}[XX']^{-1} \left(\frac{1}{n} \sum_{i=1}^{n} X_i X'_i \right) (\widehat{\beta} - \beta_0) - (\widehat{\beta} - \beta_0) \right\| \\ &= \left\| \mathbb{E}[XX']^{-1} \left(\frac{1}{n} \sum_{i=1}^{n} X_i X'_i \right) \left(\frac{1}{n} \sum_{i=1}^{n} X_i X'_i \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^{n} X_i \varepsilon_i \right) - (\widehat{\beta} - \beta_0) \right\|. \end{aligned}$$

From Equation (5.17), we obtain using the properties of the operator norm combined with Assumption 5.2(i) and Equation (5.12),

$$\begin{aligned} \left| u'\widehat{\beta} - u'\beta_0 - \overline{\xi}_n \right| &\leq \|u\| \, \left\| \mathbb{E}[XX']^{-1} \right\| \, \left\| \frac{1}{n} \sum_{i=1}^n X_i X_i' - \mathbb{E}[XX'] \right\| \, \left\| \widehat{\beta} - \beta_0 \right\| \\ &\leq \|u\| \, \frac{\left\| \widehat{\beta} - \beta_0 \right\|}{\lambda_m} \sqrt{\frac{K_{\text{reg}}}{n\delta}}. \end{aligned}$$

By Cauchy-Schwarz's inequality and Equations (5.15), we get

$$\left\|\widehat{\beta} - \beta_0\right\| \le \left\| \left(\frac{1}{n} \sum_{i=1}^n X_i X_i'\right)^{-1} \right\| \left\| \frac{1}{n} \sum_{i=1}^n X_i \varepsilon_i \right\| \le \frac{1}{\lambda_{\min}(\overline{XX'})} \sqrt{\frac{2}{n}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4}.$$
 (5.18)

Step 1c. Bound on the oracle variance $u'\mathcal{V}u$. In this step, we still reason on the event \mathcal{A} on which we prove

$$u' \mathscr{V} u \le u' \widehat{V} u + \|u\|^2 R_{n, \text{var}}(\delta)$$
(5.19)

for some $R_{n,var}(\delta)$ to be specified later and

$$\mathscr{V} := \mathbb{E}[XX']^{-1} \left(\frac{1}{n} \sum_{i=1}^n X_i X_i' \varepsilon_i^2\right) \mathbb{E}[XX']^{-1}.$$

Note that \mathscr{V} is a random variable, however it depends on the unknown $\mathbb{E}[XX']^{-1}$ and $(\varepsilon_i)_{i=1}^n$. Furthermore, $u'\mathscr{V}u = \frac{1}{n}\sum_{i=1}^n \langle u, \mathbb{E}[XX']^{-1}X_i\varepsilon_i\rangle^2$. Adding and subtracting $\hat{\varepsilon}_i$ and then expanding the square yields $\mathscr{V} = V_1 + V_2 + V_3$, where

$$V_{1} := \frac{1}{n} \sum_{i=1}^{n} \langle u, \mathbb{E}[XX']^{-1} X_{i} \widehat{\varepsilon}_{i} \rangle^{2}$$

$$V_{2} := \frac{1}{n} \sum_{i=1}^{n} \langle u, \mathbb{E}[XX']^{-1} X_{i} (\varepsilon_{i} - \widehat{\varepsilon}_{i}) \rangle^{2}$$

$$V_{3} := \frac{2}{n} \sum_{i=1}^{n} \langle u, \mathbb{E}[XX']^{-1} X_{i} \widehat{\varepsilon}_{i} \rangle \langle u, \mathbb{E}[XX']^{-1} X_{i} (\varepsilon_{i} - \widehat{\varepsilon}_{i}) \rangle.$$

By Cauchy-Schwarz and using $\varepsilon_i - \hat{\varepsilon}_i = X'_i(\hat{\beta} - \beta_0)$, as well as Equation (5.18) and Assumption 5.2 (i), we have

$$V_{2} \leq \|u\|^{2} \|\mathbb{E}[XX']^{-1}\|^{2} \|\widehat{\beta} - \beta_{0}\|^{2} \times \frac{1}{n} \sum_{i=1}^{n} \|X_{i}X'_{i}\|^{2}$$
$$\leq \|u\|^{2} \frac{1}{\lambda_{m}^{2}} \left(\frac{1}{\lambda_{\min}(\overline{XX'})} \sqrt{\frac{2}{n}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4}\right)^{2} \times \frac{1}{n} \sum_{i=1}^{n} \|X_{i}X'_{i}\|^{2}$$
$$\leq \frac{2\|u\|^{2}}{n\lambda_{m}^{2}\lambda_{\min}(\overline{XX'})^{2}} \sqrt{\frac{K_{\varepsilon}}{\delta}} \times \frac{1}{n} \sum_{i=1}^{n} \|X_{i}X'_{i}\|^{2}$$

and

$$V_{3} \leq 2 \|u\|^{2} \|\mathbb{E}[XX']^{-1}\|^{2} \|\widehat{\beta} - \beta_{0}\| \times \frac{1}{n} \sum_{i=1}^{n} \|X_{i}X'_{i}\| \|X_{i}\widehat{\varepsilon}_{i}\|$$

$$\leq \frac{2 \|u\|^{2}}{\lambda_{m}^{2}} \frac{1}{\lambda_{\min}(\overline{XX'})} \sqrt{\frac{2}{n}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4} \times \frac{1}{n} \sum_{i=1}^{n} \|X_{i}X'_{i}\| \|X_{i}\widehat{\varepsilon}_{i}\|$$

$$\leq \frac{2\sqrt{2} \|u\|^{2}}{\lambda_{min}^{2} \lambda_{\min}(\overline{XX'}) \sqrt{n}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4} \times \frac{1}{n} \sum_{i=1}^{n} \|X_{i}X'_{i}\| \|X_{i}\widehat{\varepsilon}_{i}\|.$$

Let us now focus on V_1 . We define $H_n := \mathbb{E}[XX']^{-1} - \left(\frac{1}{n}\sum_{i=1}^n X_iX'_i\right)^{-1}$. Adding and substracting $(n^{-1}\sum_{i=1}^n X_iX'_i)^{-1}$ and expanding the square, we get

$$V_1 = \frac{1}{n} \sum_{i=1}^n \langle u, \mathbb{E}[XX']^{-1} X_i \widehat{\varepsilon}_i \rangle^2 = u' \widehat{V} u + V_4 + V_5,$$

where

$$V_4 := \frac{1}{n} \sum_{i=1}^n \langle u, H_n X_i \widehat{\varepsilon}_i \rangle^2,$$

$$V_5 := \frac{2}{n} \sum_{i=1}^n \langle u, H_n X_i \widehat{\varepsilon}_i \rangle \left\langle u, \left(\frac{1}{n} \sum_{j=1}^n X_j X_j'\right)^{-1} X_i \widehat{\varepsilon}_i \right\rangle.$$

Combining the identity $A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1}$, the definition of \mathcal{A} and Equa-

tion (5.14), we get

$$||H_n||^2 \le \left\| \mathbb{E}[XX']^{-1} (\overline{XX'} - \mathbb{E}[XX']) \overline{XX'}^{-1} \right\|^2$$
$$\le \left\| \mathbb{E}[XX']^{-1} \right\|^2 \left\| \overline{XX'} - \mathbb{E}[XX'] \right\|^2 \left\| \overline{XX'}^{-1} \right\|^2$$
$$\le \frac{K_{\text{reg}}}{n\delta\lambda_m^2 \lambda_{\min}(\overline{XX'})^2}.$$

Therefore,

$$V_{4} \leq \|u\|^{2} \|H_{n}\|^{2} \times \frac{1}{n} \sum_{i=1}^{n} \|X_{i}\widehat{\varepsilon}_{i}\|^{2} \leq \frac{\|u\|^{2} K_{\text{reg}}}{n\delta\lambda_{m}^{2}\lambda_{\min}(\overline{XX'})^{2}} \times \frac{1}{n} \sum_{i=1}^{n} \|X_{i}\widehat{\varepsilon}_{i}\|^{2},$$

and

$$V_5 \le 2\|u\|^2 \|H_n\| \left\| \overline{XX'}^{-1} \right\| \times \frac{1}{n} \sum_{i=1}^n \|X_i \widehat{\varepsilon}_i\|^2 \le \frac{2\|u\|^2}{\lambda_m \lambda_{\min}(\overline{XX'})^2} \sqrt{\frac{K_{\text{reg}}}{n\delta}} \times \frac{1}{n} \sum_{i=1}^n \|X_i \widehat{\varepsilon}_i\|^2.$$

Finally, we obtain Equation (5.19), for the choice of $R_{n,var}(\delta)$ given by

$$R_{n,\text{var}}(\delta) := \frac{2}{n\lambda_m^2 \lambda_{\min}(\overline{XX'})^2} \sqrt{\frac{K_{\varepsilon}}{\delta}} \times \frac{1}{n} \sum_{i=1}^n \|X_i X_i'\|^2 + \frac{2\sqrt{2}}{\lambda_m^2 \lambda_{\min}(\overline{XX'}) \sqrt{n}} \left(\frac{K_{\varepsilon}}{\delta}\right)^{1/4} \times \frac{1}{n} \sum_{i=1}^n \|X_i X_i'\| \|X_i \widehat{\varepsilon}_i\| + \frac{K_{\text{reg}}}{n\delta \lambda_m^2 \lambda_{\min}(\overline{XX'})^2} \times \frac{1}{n} \sum_{i=1}^n \|X_i \widehat{\varepsilon}_i\|^2 + \frac{2}{\lambda_m \lambda_{\min}(\overline{XX'})^2} \sqrt{\frac{K_{\text{reg}}}{n\delta}} \times \frac{1}{n} \sum_{i=1}^n \|X_i \widehat{\varepsilon}_i\|^2.$$

Step 2. Control of the self-normalized sum $\overline{\xi}_n/\sqrt{u'\mathcal{V}u}$. Applying either Lemma 5.8 (ii) under Assumption 5.2 (iv) (in this case, $\Delta_{n,E} := \Delta_{n,E}^{\text{gen}}$), or Lemma 5.9 (ii) under Assumptions 5.2 (iv) and 5.3 (then, $\Delta_{n,E} := \Delta_{n,E}^{\text{cont}}$), we have

$$\mathbb{P}\left(\sqrt{n}\frac{|\overline{\xi}_n|}{\sqrt{u'\mathscr{V}u}} > y\right) \le 2\left(\Phi\left(-\frac{y}{\sqrt{a}}\right) + \Delta_{n,\mathrm{E}}\right) + \exp\left(-\frac{n(1-1/a)^2}{2K_{\xi}}\right)$$
(5.20)

for any given value y > 0.

Step 3. Combining the previous results. Recall we want to show Equation (5.11). We denote by \mathcal{B} the corresponding event

$$\mathcal{B} := \left\{ \frac{|u'\widehat{\beta} - u'\beta_0|}{\sqrt{u'\widehat{V}u + ||u||^2 R_{n,\operatorname{var}}(\delta)}} > Q_n^{\operatorname{Edg}} n^{-1/2} \right\}.$$

We can write by the law of total probabilities, the definition of \mathcal{A} that ensures this event arises with probability at least $1 - 2\delta$ and the fact that $\delta = \omega \alpha/2$

$$\mathbb{P}(\mathcal{B}) = \mathbb{P}(\mathcal{B} \cap \mathcal{A}) + \mathbb{P}(\mathcal{B} \cap \mathcal{A}^c) \le \mathbb{P}(\mathcal{B} \cap \mathcal{A}) + \omega\alpha.$$

Then, combining Equations (5.16) and (5.19), we get on $\mathcal{B} \cap \mathcal{A}$,

$$\frac{|u'\widehat{\beta} - u'\beta_0|}{\sqrt{u'\widehat{V}u + \|u\|^2 R_{n,\mathrm{var}}(\delta)}} \le \frac{|\overline{\xi}_n|}{\sqrt{u'\mathscr{V}u}} + \frac{R_{n,\mathrm{lin}}(\delta)/\sqrt{n}}{\sqrt{u'\widehat{V}u + \|u\|^2 R_{n,\mathrm{var}}(\delta)}}.$$

The second term in the right-hand side can be computed entirely from the data and the bounds imposed in Assumption 5.2. In contrast, $\overline{\xi}_n$ involves the unobserved random variables $(\xi_i)_{i=1}^n$. Thanks to the bound $u'\widehat{V}u + ||u||^2R_{n,\text{var}}(\delta) \geq u'\mathscr{V}u$, we recover the self-normalized sum $|\overline{\xi}_n|/\sqrt{u'\mathscr{V}u}$ that can be dealt with using the result of Step 2. More precisely, we resort to the law of total probability and apply Equation (5.20) so that, for any y > 0,

$$\mathbb{P}(\mathcal{B}) \leq \mathbb{P}\left(\frac{|\overline{\xi}_{n}|}{\sqrt{u'\mathcal{V}u}} + \frac{R_{n,\mathrm{lin}}(\delta)/\sqrt{n}}{\sqrt{u'\hat{V}u} + ||u||^{2}R_{n,\mathrm{var}}(\delta)} > Q_{n}^{\mathrm{Edg}}n^{-1/2}\right) + \omega\alpha.$$

$$\leq \mathbb{P}\left(\frac{y}{\sqrt{n}} + \frac{R_{n,\mathrm{lin}}(\delta)/\sqrt{n}}{\sqrt{u'\hat{V}u} + ||u||^{2}R_{n,\mathrm{var}}(\delta)} > Q_{n}^{\mathrm{Edg}}/\sqrt{n}\right) + \omega\alpha + \mathbb{P}\left(\sqrt{n}\frac{|\overline{\xi}_{n}|}{\sqrt{u'\mathcal{V}u}} > y\right)$$

$$\leq \mathbb{P}\left(\frac{y}{\sqrt{n}} + \frac{R_{n,\mathrm{lin}}(\delta)/\sqrt{n}}{\sqrt{u'\hat{V}u} + ||u||^{2}R_{n,\mathrm{var}}(\delta)} > Q_{n}^{\mathrm{Edg}}/\sqrt{n}\right) + \omega\alpha$$

$$+ 2\left(\Phi\left(-\frac{y}{\sqrt{a}}\right) + \Delta_{n,\mathrm{E}}\right) + \exp\left(-\frac{n(1-1/a)^{2}}{2K_{\xi}}\right).$$
(5.21)

Our goal is to choose y such that

$$2\left(\Phi\left(-\frac{y}{\sqrt{a}}\right) + \Delta_{n,\mathrm{E}}\right) + \exp\left(-\frac{n(1-1/a)^2}{2K_{\xi}}\right) = (1-\omega)\alpha.$$

Solving this equation, we find, that the solution is

$$y^* = \sqrt{a} q_{\mathcal{N}(0,1)} \left(1 - \frac{(1-\omega)\alpha - \exp\left(-n(1-1/a)^2/(2K_{\xi})\right)}{2} + \Delta_{n,\mathrm{E}} \right)$$
$$= \sqrt{a} q_{\mathcal{N}(0,1)} (1-\alpha/2 + \nu_n^{\mathrm{Edg}})$$

whenever $\nu_n^{\text{Edg}} < \alpha/2$. Therefore, we get

$$\mathbb{P}(\mathcal{B}) \leq \mathbb{P}\left(y^*/\sqrt{n} + \frac{R_{n,\mathrm{lin}}(\delta)/\sqrt{n}}{\sqrt{u'\hat{V}u + \|u\|^2 R_{n,\mathrm{var}}(\delta)}} > Q_n^{\mathrm{Edg}}/\sqrt{n}\right) + \omega\alpha + (1-\omega)\alpha = \alpha,$$

since the equality $y^*/\sqrt{n} + \frac{R_{n,\text{lin}}(\delta)/\sqrt{n}}{\sqrt{u'\widehat{V}u + \|u\|^2 R_{n,\text{var}}(\delta)}} = Q_n^{\text{Edg}}/\sqrt{n}$ holds, thus the probability in the right-hand side of the previous equation is null. \Box

5.A.2 Nonasymptotic conservativeness for the Exponential regime

Lemma 5.2. For every $\alpha \in (0,1)$, a > 1, $n \ge 2$ and $\omega \in (0,1)$, if $n\omega > 2K_{\text{reg}}/(\alpha \lambda_m^2)$ and $\nu_n^{\text{Exp}} < \alpha/2$, we have

$$\mathbb{P}\left(|u'\widehat{\beta} - u'\beta_0| \le n^{-1/2}Q_n^{\operatorname{Exp}}\sqrt{u'\widehat{V}u + ||u||^2 R_{n,\operatorname{var}}(\omega\alpha/2)}\right) \ge 1 - \alpha.$$

Proof of Lemma 5.2. This proof is very similar to the one given in Section 5.A.1. Steps 1.a-b-c are left unchanged. In the second step, we use an exponential bound to control directly the self-normalized sum instead of using Lemmas 5.8 (ii) or 5.9 (ii) that involve Edgeworth expansions.

By (Bertail et al., 2008, Equation (6)), we get for every t > 0,

$$\mathbb{P}\left(n\frac{|\bar{\xi}_n|^2}{u'\mathscr{V}u} > t\right) \le 2e^{1-t/(2(1+a))} + e^{-n(1-1/a)^2/(2K_{\xi})},$$

and therefore, by substituting $y = \sqrt{t}$, we get for every y > 0

$$\mathbb{P}\left(\sqrt{n}\frac{|\bar{\xi}_n|}{\sqrt{u'\mathscr{V}u}} > y\right) \le 2e^{1-y^2/(2(1+a))} + e^{-n(1-1/a)^2/(2K_{\xi})} = 2e^{1-y^2/(2(1+a))} + 2\nu_n^{\mathrm{Exp}} - \omega\alpha$$

by definition of ν_n^{Exp} . Now, we look for y such that $2e^{1-y^2/(2(1+a))} + 2\nu_n^{\text{Exp}} = \alpha$. We find, whenever $\nu_n^{\text{Exp}} < \alpha/2$,

$$y^* = \sqrt{2(1+a)\ln\left(\frac{e}{\alpha/2 - \nu_n^{\text{Exp}}}\right)}.$$

We denote by \mathcal{B} the corresponding event

$$\mathcal{B} := \left\{ \frac{|u'\widehat{\beta} - u'\beta_0|}{\sqrt{u'\widehat{V}u + ||u||^2 R_{n,\mathrm{var}}(\delta)}} > Q_n^{\mathrm{Exp}}/\sqrt{n} \right\}.$$

The reasoning concluding the proof of Lemma 5.1 applies here as well. Therefore, we get

$$\mathbb{P}(\mathcal{B}) \leq \mathbb{P}\left(y^* / \sqrt{n} + \frac{R_{n,\mathrm{lin}}(\delta) / \sqrt{n}}{\sqrt{u'\hat{V}u + \|u\|^2 R_{n,\mathrm{var}}(\delta)}} > Q_n^{\mathrm{Exp}} / \sqrt{n}\right) + \omega\alpha + (1-\omega)\alpha = \alpha,$$

since the equality $y^*/\sqrt{n} + \frac{R_{n,\text{lin}}(\delta)/\sqrt{n}}{\sqrt{u'\hat{V}u + \|u\|^2 R_{n,\text{var}}(\delta)}} = Q_n^{\text{Exp}}/\sqrt{n}$ holds, thus the probability in the right-hand side of the previous equation is null. \Box

5.A.3 Proof of Theorem 5.1

For fixed α , n, ω and a, we remark that exactly one case out of the three that intervene in definition of $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$ arises. Furthermore, the conditions defining these cases are deterministic. As a consequence, we can consider each case separately and check that the coverage of $u'\beta_0$ is at least $1-\alpha$.

If
$$n \leq 2K_{\text{reg}}/(\omega \alpha \lambda_m^2)$$
 or $\nu_n^{\text{Exp}} \geq \alpha/2$, $\operatorname{CI}_u^{\text{Fin}}(1-\alpha, n) = \mathbb{R}$ so that $\mathbb{P}\left(u'\beta_0 \in \operatorname{CI}_u^{\text{Fin}}(1-\alpha, n)\right) = 1 > 1-\alpha.$

Otherwise, if $\nu_n^{\text{Edg}} \ge \alpha/2$, $\operatorname{CI}_u^{\text{Fin}}(1-\alpha,n) = \operatorname{CI}_u^{\text{Exp}}(1-\alpha,n)$. In this case, Lemma 5.2 ensures that

$$\mathbb{P}\left(u'\beta_0 \in \mathrm{CI}_u^{\mathrm{Fin}}(1-\alpha,n)\right) \ge 1-\alpha$$

If not, that is, $\nu_n^{\text{Edg}} < \alpha/2$, we are in the Edg regime and $\operatorname{CI}_u^{\text{Fin}}(1-\alpha, n) = \operatorname{CI}_u^{\text{Edg}}(1-\alpha, n)$. In this case, Lemma 5.1 guarantees the result. \Box
5.A.4 Proof of Theorem 5.2

The proof is divided in two main steps. First, we show that the relevant regime asymptotically is the Edgeworth one. Second, we show that

$$\mathbb{P}\left(u'\beta_0 \in \mathrm{CI}_u^{\mathrm{Edg}}(1-\alpha,n)\right) \to 1-\alpha.$$
(5.22)

Step 1. We start by proving that ν_n^{Edg} tends to 0. Since $a_n = 1 + b_n$ with $b_n = o(1)$ and $b_n \sqrt{n} \to +\infty$, we have

$$\sqrt{n}\left(1-\frac{1}{a_n}\right) \to +\infty.$$

Furthermore, by assumption, $\omega_n \to 0$. Combining those two limits, we obtain

$$\nu_n^{\rm Exp} = \frac{\omega_n \alpha + \exp\left(-n(1-1/a_n)^2/(2K_{\xi})\right)}{2} \to 0.$$
 (5.23)

Besides, the discussion following the definition of $\Delta_{n,E}$ in the body of the article implies that under Assumption 5.2 (iv), $\Delta_{n,E} = \Delta_{n,E}^{\text{gen}} = o(1)$ (respectively, under Assumptions 5.2 (iv) and 5.3, $\Delta_{n,E} = \Delta_{n,E}^{\text{cont}} = o(1)$). From this and (5.23), we conclude that $\nu_n^{\text{Edg}} = \nu_n^{\text{Exp}} + \Delta_{n,E} = o(1)$.

This implies that for n large enough (depending on α and the bounds of Assumption 5.2), $\nu_n^{\text{Edg}} < \alpha/2$ and therefore we can only be in the "Edg" regime, meaning that

$$\mathbb{P}\left(u'\beta_0 \in \mathrm{CI}_u^{\mathrm{Fin}}(1-\alpha,n)\right) = \mathbb{P}\left(u'\beta_0 \in \mathrm{CI}_u^{\mathrm{Edg}}(1-\alpha,n)\right)$$

We thus focus below on the Edgeworth regime.

Step 2. In order to prove (5.22), we show that

$$\frac{q_{\mathcal{N}(0,1)}(1-\alpha/2)}{Q_n^{\mathrm{Edg}}} \frac{\sqrt{n}u'(\hat{\beta}-\beta_0)}{\sqrt{u'\hat{V}u+\|u\|^2 R_{n,\mathrm{var}}(\omega_n\alpha/2)}} \xrightarrow{d} \mathcal{N}(0,1).$$
(5.24)

First, we remark that under the assumptions of the theorem, we are in the parameter set Θ^{Asymp} . Therefore, classical results (see van der Vaart (2000), Chapters 2, 3 and 5 for instance) allow us to claim

$$\frac{\sqrt{n}u'(\hat{\beta} - \beta_0)}{\sqrt{u'\hat{V}u}} \xrightarrow{d} \mathcal{N}(0, 1).$$
(5.25)

Under the assumptions $\sqrt{n}\omega_n^{3/4} \to +\infty$, $\mathbb{E}[||X\varepsilon||^2] < +\infty$ and $\mathbb{E}[||X||^4] < +\infty$, we have

 $R_{n,\text{var}}(\omega_n \alpha/2) \xrightarrow{P} 0 \text{ and } R_{n,\text{lin}}(\omega_n \alpha/2) \xrightarrow{P} 0.$

In addition, we also have $u'\hat{V}u \xrightarrow{P} u'Vu > 0$ since we place ourselves in Θ^{Asymp} . Those three results imply that $\nu_n^{\text{Approx}} \xrightarrow{P} 0$.

By the convergence of $R_{n,var}(\omega_n \alpha/2)$ and the Continuous Mapping Theorem (CMT), we also have

$$\frac{\sqrt{u'\widehat{V}u + \|u\|^2 R_{n,\text{var}}(\omega_n \alpha/2)}}{\sqrt{u'\widehat{V}u}} \xrightarrow{P} 1.$$
(5.26)

From Step 1, we also know that $\nu_n^{\text{Edg}} \to 0$. Combined with the assumption $a_n = 1 + o(1)$, we obtain

$$Q_n^{\text{Edg}} \xrightarrow{P} q_{\mathcal{N}(0,1)}(1 - \alpha/2).$$
(5.27)

Equations (5.25), (5.26) and (5.27) and Slutsky's lemma gives the desired result of Equation (5.24). The latter implies that the probability

$$\mathbb{P}\left(-q_{\mathcal{N}(0,1)}(1-\alpha/2) \le \frac{q_{\mathcal{N}(0,1)}(1-\alpha/2)}{Q_n^{\text{Edg}}} \frac{\sqrt{n}u'(\hat{\beta}-\beta_0)}{\sqrt{u'\hat{V}u + \|u\|^2 R_{n,\text{var}}(\omega_n\alpha/2)}} \le q_{\mathcal{N}(0,1)}(1-\alpha/2)\right)$$

converges to $1 - \alpha$. Finally, the event considered in this probability is equivalent to

$$\left\{-\frac{Q_n^{\text{Edg}}}{\sqrt{n}} \le \frac{u'(\hat{\beta} - \beta_0)}{\sqrt{u'\hat{V}u + \|u\|^2 R_{n,\text{var}}(\omega_n \alpha/2)}} \le \frac{Q_n^{\text{Edg}}}{\sqrt{n}}\right\},\,$$

also equivalent to the event $\{u'\beta_0 \in \operatorname{CI}_u^{\operatorname{Edg}}(1-\alpha,n)\}$, which is what we want given Step 1. \Box

Appendix 5.B Proof of results in Section 5.6

5.B.1 Nonasymptotic conservativeness for the Edgeworth and Exponential regimes

Lemma 5.3. For every $\alpha \in (0,1)$, a > 1, $n \ge 1$ and $\omega \in (0,1)$, if $n\omega > 2K_{\text{reg}}/(\alpha \lambda_m^2)$ and $\nu_n^{\text{Edg}} < \alpha/2$, we have

$$\mathbb{P}\left(|\mathbf{e}_{2}'\widehat{\psi}(\beta_{0,2})| \leq \frac{Q_{n}^{\mathrm{Edg}}(\beta_{0,2})}{\sqrt{n}}\sqrt{\mathbf{e}_{2}'\widehat{V}(\beta_{0,2})\mathbf{e}_{2} + R_{n,\mathrm{var}}(\omega\alpha/2,\beta_{0,2})}\right) \geq 1 - \alpha.$$

Proof of Lemma 5.3. We recall that $\psi(\beta_{0,2}) = 0$. Hence, the stated lemma is equivalent to

$$\mathbb{P}\Big(|\mathbf{e}_{2}'\hat{\psi}(\beta_{0,2}) - \mathbf{e}_{2}'\psi(\beta_{0,2})| \le \frac{Q_{n}^{\mathrm{Edg}}(\beta_{0,2})}{\sqrt{n}}\sqrt{\mathbf{e}_{2}'\hat{V}(\beta_{0,2})\mathbf{e}_{2} + R_{n,\mathrm{var}}(\omega\alpha/2,\beta_{0,2})}\Big) \ge 1 - \alpha.$$

Since $\mathbf{e}'_2 \hat{\psi}(\beta_{0,2})$ (resp. $\mathbf{e}'_2 \psi(\beta_{0,2})$) is the OLS estimator in a regression of $Y - \beta_{0,2}D$ on W (resp. its theoretical counterpart), the proof of Lemma 5.1 with $u = \mathbf{e}_2$ can be replicated under Assumption 5.5 up to minor notational changes (see the definition of $Q_n^{\text{Edg}}(\beta_{0,2})$, $\hat{V}(\beta_{0,2})$ or $R_{n,\text{var}}(\omega\alpha/2,\beta_{0,2})$ in Section 5.6). \Box

Lemma 5.4. For every $\alpha \in (0, 1)$, a > 1, $n \ge 2$ and $\omega \in (0, 1)$, if $n\omega > 2K_{\text{reg}}/(\alpha \lambda_m^2)$ and $\nu_n^{\text{Exp}} < \alpha/2$, we have

$$\mathbb{P}\Big(|\mathbf{e}_{2}'\hat{\psi}(\beta_{0,2})| \le n^{-1/2}Q_{n}^{\mathrm{Exp}}(\beta_{0,2})\sqrt{\mathbf{e}_{2}'\hat{V}(\beta_{0,2})\mathbf{e}_{2} + R_{n,\mathrm{var}}(\omega\alpha/2,\beta_{0,2})}\Big) \ge 1 - \alpha$$

Proof of Lemma 5.4. The remarks put forward in the proof of Lemma 5.3 remain valid and allow us to conclude that the proof of Lemma 5.2 can be replicated under Assumption 5.5 up to minor notational changes. \Box

5.B.2 Proof of Theorem 5.3

For fixed α , n, ω and a, we remark that exactly one case out of the three that intervene in definition of $CS^{Fin}(1-\alpha, n)$ arises. Furthermore, the conditions defining these cases are deterministic. As a consequence, we can consider each case separately and check that the coverage of $\beta_{0,2}$ is at least $1-\alpha$.

If
$$n \leq 2K_{\text{reg}}/(\omega \alpha \lambda_m^2)$$
 or $\nu_n^{\text{Exp}} \geq \alpha/2$, $\text{CS}^{\text{Fin}}(1-\alpha, n) = \mathbb{R}$ so that
 $\mathbb{P}\Big(\beta_{0,2} \in \text{CS}^{\text{Fin}}(1-\alpha, n)\Big) = 1 > 1-\alpha.$

Otherwise, if $\nu_n^{\text{Edg}} \ge \alpha/2$, $\text{CS}^{\text{Fin}}(1-\alpha, n) = \text{CS}^{\text{Exp}}(1-\alpha, n)$. Given the definition of $\text{CS}^{\text{Exp}}(1-\alpha, n)$, we can write in that case

$$\mathbb{P}\Big(\beta_{0,2} \in \mathrm{CS}^{\mathrm{Fin}}(1-\alpha,n)\Big) = \mathbb{P}\Big(\beta_{0,2} \in \mathrm{CS}^{\mathrm{Exp}}(1-\alpha,n)\Big)$$
$$= \mathbb{P}\Big(|\mathbf{e}_{2}'\hat{\psi}(\beta_{0,2})| \leq \frac{Q_{n}^{\mathrm{Exp}}(\beta_{0,2})}{\sqrt{n}}\sqrt{\mathbf{e}_{2}'\hat{V}(\beta_{0,2})\mathbf{e}_{2} + R_{n,\mathrm{var}}(\omega\alpha/2,\beta_{0,2})}\Big)$$

which is at least $1 - \alpha$ by Lemma 5.4.

In the final scenario, that is when $n > 2K_{\text{reg}}/(\omega \alpha \lambda_m^2)$, $\nu_n^{\text{Exp}} < \alpha/2$ and $\nu_n^{\text{Edg}} < \alpha/2$, we are in the Edg regime and $\text{CS}^{\text{Fin}}(1-\alpha,n) = \text{CS}^{\text{Edg}}(1-\alpha,n)$. In this case, Lemma 5.3 guarantees the result. \Box

5.B.3 Proof of Theorem 5.4

Step 1 in the proof of Theorem 5.2 can be directly adapted to conclude that for n large enough (depending on α and the bounds of Assumption 5.5), $\nu_n^{\text{Edg}} < \alpha/2$. Therefore we can only be in the "Edg" regime for n large enough, meaning that

$$\mathbb{P}\left(\beta_{0,2} \in \mathrm{CS}^{\mathrm{Fin}}(1-\alpha,n)\right) = \mathbb{P}\left(\beta_{0,2} \in \mathrm{CS}^{\mathrm{Edg}}(1-\alpha,n)\right).$$

We thus focus below on the Edgeworth regime and seek to prove

$$\mathbb{P}\left(\beta_{0,2} \in \mathrm{CS}^{\mathrm{Edg}}(1-\alpha, n)\right) \to 1-\alpha.$$
(5.28)

In order to prove (5.28), we can follow the proof of Theorem 5.2 to show that

$$\frac{q_{\mathcal{N}(0,1)}(1-\alpha/2)}{Q_n^{\mathrm{Edg}}(\beta_{0,2})} \frac{\sqrt{n}\mathbf{e}_2'\psi(\beta_{0,2})}{\sqrt{\mathbf{e}_2'\hat{V}(\beta_{0,2})\mathbf{e}_2 + R_{n,\mathrm{var}}(\omega_n\alpha/2,\beta_{0,2})}} \xrightarrow{d} \mathcal{N}(0,1).$$
(5.29)

The latter implies that the probability

$$\mathbb{P}\left(-q_{\mathcal{N}(0,1)}(1-\alpha/2) \le \frac{q_{\mathcal{N}(0,1)}(1-\alpha/2)}{Q_n^{\mathrm{Edg}}(\beta_{0,2})} \frac{\sqrt{n}\mathbf{e}_2'\hat{\psi}(\beta_{0,2})}{\sqrt{\mathbf{e}_2'\hat{V}(\beta_{0,2})\mathbf{e}_2 + R_{n,\mathrm{var}}(\omega_n\alpha/2,\beta_{0,2})}} \le q_{\mathcal{N}(0,1)}(1-\alpha/2)\right)$$

converges to $1 - \alpha$. Finally, the event considered in this probability is equivalent to

$$\left\{-\frac{Q_n^{\mathrm{Edg}}(\beta_{0,2})}{\sqrt{n}} \le \frac{\mathbf{e}_2'\hat{\psi}(\beta_{0,2})}{\sqrt{\mathbf{e}_2'\hat{V}(\beta_{0,2})\mathbf{e}_2 + R_{n,\mathrm{var}}(\omega_n\alpha/2,\beta_{0,2})}} \le \frac{Q_n^{\mathrm{Edg}}(\beta_{0,2})}{\sqrt{n}}\right\},$$

also equivalent to the event $\{\beta_{0,2} \in \mathrm{CS}^{\mathrm{Edg}}(1-\alpha,n)\}$, which is what we want. \Box

Appendix 5.C Additional lemmas

Lemma 5.5. Let $n \ge 1$ and $(A_i)_{i=1}^n$ be an *i.i.d.* sequence of random square matrices of dimension d with finite second moment. For every $\delta \in (0, 1)$,

$$\mathbb{P}\left(\left|\left|\frac{1}{n}\sum_{i=1}^{n}A_{i}-\mathbb{E}[A]\right|\right|>\sqrt{\frac{B}{n\delta}}\right)<\delta,$$

where $B := \mathbb{E}\left[\left|\left|vec(A - \mathbb{E}[A])\right|\right|^2\right]$.

Proof of Lemma 5.5: By Markov's inequality and the inequality $||M||^2 \leq \sum_{1 \leq l, l' \leq d} (M^{(l,l')})^2$ valid for every square matrix M, we get for every t > 0

$$\mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^{n}A_{i}-\mathbb{E}[A]\right\|>t\right)\leq\frac{1}{t^{2}}\mathbb{E}\left[\left\|\frac{1}{n}\sum_{i=1}^{n}A_{i}-\mathbb{E}[A]\right\|^{2}\right]$$
$$\leq\frac{1}{t^{2}}\mathbb{E}\left[\sum_{1\leq l,l'\leq d}\left(\frac{1}{n}\sum_{i=1}^{n}A_{i}^{(l,l')}-\mathbb{E}\left[A^{(l,l')}\right]\right)^{2}\right].$$

Since $(A_i)_{i=1}^n$ is an *i.i.d.* sequence, we can write

$$\mathbb{E}\left[\sum_{1\leq l,l'\leq d} \left(\frac{1}{n}\sum_{i=1}^{n}A_{i}^{(l,l')} - \mathbb{E}\left[A^{(l,l')}\right]\right)^{2}\right] = \frac{1}{n}\sum_{1\leq l,l'\leq d}\mathbb{E}\left[\left(A^{(l,l')} - \mathbb{E}\left[A^{(l,l')}\right]\right)^{2}\right]$$
$$= \frac{1}{n}\mathbb{E}\left[\left|\left|vec(A - \mathbb{E}[A])\right|\right|^{2}\right].$$

Choosing $t = \sqrt{\frac{B}{n\delta}}$ concludes the proof. \Box

For reader's convenience, we recall the following eigenvalue stability lemma, which is a corollary of Weyl's inequality $\lambda_{min}(A+B) \leq \lambda_{min}(A) + \lambda_{max}(B)$ for real symmetric matrices, see for example (Hogben, 2006, Section 8.2).

Lemma 5.6. For M_1 and M_2 be two symmetric matrices of dimension d, we have $|\lambda_{min}(M_1) - \lambda_{min}(M_2)| \leq ||M_1 - M_2||.$

Proof of Lemma 5.6: If $\lambda_{min}(M_1) = \lambda_{min}(M_2)$, the results follows directly. Without loss of generality, $\lambda_{min}(M_1) > \lambda_{min}(M_2)$. Applying Weyl's inequality with $A := M_2$ and $B := M_1 - M_2$ we get $0 < \lambda_{min}(M_1) - \lambda_{min}(M_2) \le \lambda_{max}(M_1 - M_2) \le ||M_1 - M_2||$. \Box

Lemma 5.7. Let $n \ge 1$ and $(X_i)_{i=1}^n$ be an *i.i.d.* sequence of random vectors of dimension p and $(\varepsilon_i)_{i=1}^n$ an *i.i.d.* sequence of random variables such that $\mathbb{E}[X\varepsilon] = 0$ and $\mathbb{E}[||X\varepsilon||^4] < +\infty$. For every $\delta \in (0, 1)$,

$$\mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^{n}X_{i}\varepsilon_{i}\right\| \leq \sqrt{\frac{2}{n}}\left(\frac{\mathbb{E}\left[\left\|X\varepsilon\right\|^{4}\right]}{\delta}\right)^{1/4}\right) \geq 1-\delta.$$

Proof of Lemma 5.7: By Markov's inequality, we can write

$$\mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^{n}X_{i}\varepsilon_{i}\right\| > t\right) \leq \frac{1}{(tn)^{4}}\mathbb{E}\left[\left\|\sum_{i=1}^{n}X_{i}\varepsilon_{i}\right\|^{4}\right].$$
(5.30)

We now focus on $\mathbb{E}\left[\left|\left|\sum_{i=1}^{n} X_i \varepsilon_i\right|\right|^4\right]$. We have

$$\mathbb{E}\left[\left\|\sum_{i=1}^{n} X_{i}\varepsilon_{i}\right\|^{4}\right] = \sum_{1 \leq l,l' \leq d} \sum_{1 \leq i,j,i',j' \leq n} \mathbb{E}\left[X_{il}\varepsilon_{i}X_{jl}\varepsilon_{j}X_{i'l'}\varepsilon_{i'}X_{j'l'}\varepsilon_{j'}\right]$$
(5.31)

The sequence $(X_i \varepsilon_i)_{i=1}^n$ is *i.i.d.* and centered. Consequently,

$$\mathbb{E}\Big[X_{il}\varepsilon_i X_{jl}\varepsilon_j X_{i'l'}\varepsilon_{i'} X_{j'l'}\varepsilon_{j'}\Big]$$

= $\mathbb{E}\left[X_l^2 X_{l'}^2 \varepsilon^4\right] \mathbb{1}_{\{i=j=i'=j'\}} + \mathbb{E}\left[X_l^2 \varepsilon^2\right] \mathbb{E}\left[X_{l'}^2 \varepsilon^2\right] \mathbb{1}_{\{i=j,i'=j',i\neq i'\}}$
+ $\mathbb{E}\left[X_l X_{l'} \varepsilon^2\right]^2 \left\{\mathbb{1}_{\{i=i',j=j',i\neq j\}} + \mathbb{1}_{\{i=j',j=i',i\neq j\}}\right\}.$

Combining this with (5.31), we get (using Jensen's inequality on the last line)

$$\mathbb{E}\left[\left\|\sum_{i=1}^{n} X_{i}\varepsilon_{i}\right\|^{4}\right] = \sum_{1 \leq l, l' \leq d} \left\{ n\mathbb{E}\left[X_{l}^{2}X_{l'}^{2}\varepsilon^{4}\right] + n(n-1)\mathbb{E}\left[X_{l}^{2}\varepsilon^{2}\right]\mathbb{E}\left[X_{l'}^{2}\varepsilon^{2}\right] + 2n(n-1)\mathbb{E}\left[X_{l}X_{l'}\varepsilon^{2}\right]^{2} \right\} \\
= n\mathbb{E}\left[\left||X\varepsilon||^{4}\right] + n(n-1)\mathbb{E}\left[\left||X\varepsilon||^{2}\right]^{2} + 2n(n-1)\sum_{1 \leq l, l' \leq d} \mathbb{E}\left[X_{l}X_{l'}\varepsilon^{2}\right]^{2} \\
\leq 4n^{2}\mathbb{E}\left[\left||X\varepsilon||^{4}\right].$$
(5.32)

We plug (5.32) back in (5.30) and choose $t = \sqrt{\frac{2}{n}} \left(\frac{\mathbb{E}[||X_{\varepsilon}||^4]}{\delta}\right)^{1/4}$ to conclude. \Box

Lemma 5.8 (Edgeworth expansion, general case). Let $n \ge 1$ and $(Z_i)_{i=1}^n$ be an *i.i.d.* sequence of random variables with $\mathbb{E}[Z] = 0$, $\mathbb{V}(Z) = \sigma^2$ and $\mathbb{E}[Z^4]/\sigma^4 \le K$. Let $\overline{Z}_n := n^{-1} \sum_{i=1}^n Z_i$ and $\widehat{\sigma}_0^2 := \frac{1}{n} \sum_{i=1}^n Z_i^2$. For every x > 0 and a > 1

(i)
$$\mathbb{P}\left(\sqrt{n}\left|\overline{Z}_{n}\right|/\sigma > x\right) \leq 2\left\{\Phi(-x) + \Delta_{n,E}^{\text{gen}}\right\}$$

(ii) $\mathbb{P}\left(\sqrt{n}\left|\overline{Z}_{n}\right|/\widehat{\sigma}_{0} > x\right) \leq 2\left\{\Phi(-x/\sqrt{a}) + \Delta_{n,E}^{\text{gen}}\right\} + \exp\left(-\frac{n(1-1/a)^{2}}{2K}\right)$

where $\Delta_{n,E}^{\text{gen}} := \frac{0.2(K^{3/4}+1)}{\sqrt{n}} + \frac{0.2K+0.12K^{3/2}}{n} + r_n^{\text{gen}}$ and $r_n^{\text{gen}} = O(n^{-5/4})$ is an explicit term given in Section 5.D.1 which depends on K and n only.

Proof of Lemma 5.8:

Let
$$\lambda_3 := \mathbb{E}[Z^3]/\sigma^3$$
, $K_p := \mathbb{E}[|Z|^p]/\sigma^p$ and $\operatorname{Edg}_n(x) := \frac{\lambda_3}{6\sqrt{n}}(1-x^2)\varphi(x)$. Combining

the parity of Edg_n with the fact that $\Phi(x) = 1 - \Phi(-x)$, we obtain

$$\mathbb{P}\left(\frac{\sqrt{n}\left|\overline{Z}_{n}\right|}{\sigma} > x\right) \leq 1 - \mathbb{P}\left(\frac{\sqrt{n}\overline{Z}_{n}}{\sigma} \leq x\right) + \mathbb{P}\left(\frac{\sqrt{n}\overline{Z}_{n}}{\sigma} \leq -x\right) \\
\leq 1 - \left(\mathbb{P}\left(\frac{\sqrt{n}\overline{Z}_{n}}{\sigma} \leq x\right) - \Phi(x) - \operatorname{Edg}_{n}(x)\right) \\
+ \left(\mathbb{P}\left(\frac{\sqrt{n}\overline{Z}_{n}}{\sigma} \leq -x\right) - \Phi(-x) - \operatorname{Edg}_{n}(-x)\right) \\
- \Phi(x) - \operatorname{Edg}_{n}(x) + \Phi(-x) + \operatorname{Edg}_{n}(-x) \\
\leq 2\Phi(-x) + \left|\mathbb{P}\left(\frac{\sqrt{n}\overline{Z}_{n}}{\sigma} \leq x\right) - \Phi(x) - \operatorname{Edg}_{n}(x)\right| \\
+ \left|\mathbb{P}\left(\frac{\sqrt{n}\overline{Z}_{n}}{\sigma} \leq -x\right) - \Phi(-x) - \operatorname{Edg}_{n}(-x)\right|.$$
(5.33)

As a consequence of Theorem 2.3 in Derumigny et al. (2021) in the *i.i.d.* case, we have

$$\sup_{x \in \mathbb{R}} \left| \mathbb{P}\left(\frac{\sqrt{n}\overline{Z}_n}{\sigma} \le x \right) - \Phi(x) - \operatorname{Edg}_n(x) \right| \le \frac{0.1995(K_3 + 1)}{\sqrt{n}} + \frac{0.031(K_3 + 1)^2 + 0.195K_4 + 0.054|\lambda_3|(K_3 + 1) + 0.038\lambda_3^2}{n} + r_n^{\operatorname{gen}},$$

with $r_n^{\text{gen}} = r_n^{\text{gen}}(K)$ is a remainder term given in Section 5.D.1, which bounds above the remainder $r_{1,n}$ defined in Equation (22) in Derumigny et al. (2021).

An upper bound K on K_4 is enough to control the right-hand side of the previous equation through (i) a convexity argument: $K_3 \leq K_4^{3/4}$, $K_4 \geq 1$; and (ii) Theorem 1 of Pinelis (2011*a*): $|\lambda_3| \leq 0.621K_3$.

Consequently, with $r_n \leq R_n := R_n(K)$ which is defined in Section 5.D, we obtain the following upper bound that relies only on K, an upper bound on K_4 .

$$\begin{split} \sup_{x \in \mathbb{R}} \left| \mathbb{P}\left(\frac{\sqrt{nZ_n}}{\sigma} \le x \right) - \Phi(x) - \mathrm{Edg}_n(x) \right| \\ \le \frac{0.1995(K^{3/4} + 1)}{\sqrt{n}} + \frac{0.031(K^{3/4} + 1)^2 + 0.195K + 0.0336K^{3/4}(K^{3/4} + 1) + 0.0147K^{3/2}}{n} + r_n^{\mathrm{gen}} \\ \le \frac{0.1995(K^{3/4} + 1)}{\sqrt{n}} + \frac{0.206K^{3/2} + 0.195K}{n} + r_n^{\mathrm{gen}}, \end{split}$$
(5.34)

where $r_n^{\text{gen}} = O(n^{-5/4})$. To prove this, we remark that, since the constants in the secondorder term divided by n are less important, we use the lower bound $1 \le K^{3/4}$. The numerical constants are computed as follows: $0.054 \times 0.621 < 0.0336$, $0.038 \times 0.621^2 < 0.0147$ (first inequality), and $0.031 \times 2^2 + 0.0336 \times 2 + 0.0147 < 0.206$ (second inequality).

Letting $\Delta_{n,E}^{\text{gen}} := \frac{0.2(K^{3/4}+1)}{\sqrt{n}} + \frac{0.206K^{3/2}+0.195K}{n} + R_n$ and combining (5.33) and (5.34) is enough to obtain the first result.

For the second result, we use Lemma 3 in Bertail et al. (2008) which allows us to write for every a > 1

$$\mathbb{P}\left(\frac{\widehat{\sigma}_0^2}{\sigma^2} < \frac{1}{a}\right) \le \exp\left(-\frac{n(1-1/a)^2}{2K}\right).$$
(5.35)

Combining Lemma 5.8.(i) and (5.35), we can claim for every x > 0 and a > 1

$$\mathbb{P}\left(\sqrt{n}\left|\overline{Z}_{n}\right|/\widehat{\sigma}_{0} > x\right) \leq \mathbb{P}\left(\sqrt{n}\left|\overline{Z}_{n}\right|/\sigma > x/\sqrt{a}\right) + \mathbb{P}\left(\frac{\widehat{\sigma}_{0}^{2}}{\sigma^{2}} < \frac{1}{a}\right)$$
$$\leq 2\left\{\Phi(-x/\sqrt{a}) + \Delta_{n,\mathrm{E}}\right\} + \exp\left(-\frac{n(1-1/a)^{2}}{2K}\right),$$

hence result (ii). \Box

Lemma 5.9 (Edgeworth expansion, continuous case). Let $n \ge 1$ and $(Z_i)_{i=1}^n$ be an *i.i.d.* sequence of random variables such that Z has an absolutely continuous component with respect to Lebesgue's measure, $\mathbb{E}[Z] = 0$, $\mathbb{V}(Z) = \sigma^2$, and $\mathbb{E}[Z^4] / \sigma^4 \le K$. Let $\overline{Z}_n := n^{-1} \sum_{i=1}^n Z_i$ and $\widehat{\sigma}_0^2 := \frac{1}{n} \sum_{i=1}^n Z_i^2$. For every x > 0 and a > 1

(i)
$$\mathbb{P}\left(\sqrt{n}\left|\overline{Z}_{n}\right|/\sigma > x\right) \leq 2\left\{\Phi(-x) + \Delta_{n,E}^{\text{cont}}\right\}$$

(ii) $\mathbb{P}\left(\sqrt{n}\left|\overline{Z}_{n}\right|/\widehat{\sigma}_{0} > x\right) \leq 2\left\{\Phi(-x/\sqrt{a}) + \Delta_{n,E}^{\text{cont}}\right\} + \exp\left(-\frac{n(1-1/a)^{2}}{2K}\right)$

where $\Delta_{n,E}^{\text{cont}} := \frac{0.195 K + 0.0147 K^{3/2}}{n} + r_n^{\text{cont}}$ and $r_n^{\text{cont}} = O(n^{-5/4})$ is an explicit term given in Section 5.D.2 which depends on κ (defined in Assumption 5.3), K and n only.

Proof of Lemma 5.9:

The structure of the proof follows the one of Lemma 5.8. The sole difference is the upper bound on

$$\left| \mathbb{P}\left(\frac{\sqrt{n}\overline{Z}_n}{\sigma} \le -x \right) - \Phi(-x) - \operatorname{Edg}_n(-x) \right|,$$

which comes from Corollary 3.4 in Derumigny et al. (2021). Here, the fact that the distribution of the observations does not vary with n and the upper bound $\widetilde{K}_{3,n} \leq K_{3,n} + 1 \leq K^{3/4} + 1$ implies that the term denoted by $\kappa_n(\widetilde{K}_{3,n})$ in Derumigny et al. (2021) is a positive constant strictly smaller than 1. Consequently, the second term in Corollary 3.4 in Derumigny et al. (2021) is $O(n^{-5/4})$ and we obtain

$$\left| \mathbb{P}\left(\frac{\sqrt{nZ_n}}{\sigma} \le -x\right) - \Phi(-x) - \operatorname{Edg}_n(-x) \right| \le \frac{0.195 \, K_{4,n} + 0.038 \, \lambda_{3,n}^2}{n} + r_n^{\operatorname{cont}}$$

where $r_n^{\text{cont}} = O(n^{-5/4})$ contains upper bounds on the second term and remainder in Corollary 3.4 in Derumigny et al. (2021) and is detailed in Section 5.D.2.

Result (i) given in the lemma follows from the inequalities $\lambda_{3,n} \leq 0.621 K_{3,n}$ and $K_{3,n} \leq K_{4,n}^{3/4} \leq K^{3/4}$ (0.038 × 0.621² < 0.0147). Result (ii) follows from the same reasoning as the result (ii) of Lemma 5.8. \Box

Appendix 5.D Definition of the remainders of the bound on Edgeworth expansion

5.D.1 Definition of r_n^{gen}

We define r_n^{gen} by:

$$r_n^{\text{gen}} := \frac{(14.1961 + 67.0415)(K^{3/4} + 1)^4}{2^4 \pi^4 n^2} + \frac{0.621K^{3/4} \exp(-2n^2/(K^{3/4} + 1)^4)}{3\pi\sqrt{n}} \\ + \frac{K^{3/4}}{3\sqrt{n}} J_3(3, \sqrt{0.3}(n/K)^{1/4}, T/\pi, K^{3/4} + 1, T, n) + I_{4,3}(T) \\ + \frac{0.504(e_{2,n} - 1.0067)K^{3/2}}{36n} + \frac{1.0253}{\pi} \int_0^{\sqrt{0.2}(n/K)^{1/4}} u e^{-u^2/2} R_{2,n}(u) du$$
(5.36)

where $T := 2\pi \sqrt{n} / \widetilde{K}_{3,n}$,

$$J_3(p, v, w, K, T, n) := \frac{1}{T} \int_v^w |\Psi(u/T)| u^p \exp\left(-\frac{u^2}{2}\left(1 - \frac{2\chi_1|u|K}{\sqrt{n}} - \frac{1}{n}\right)\right) du,$$

 χ_1 denotes the constant $\chi_1 := \sup_{x>0} x^{-3} |\cos(x) - 1 + x^2/2| \approx 0.099162$ (Shevtsova, 2010), and

$$\begin{split} R_{2,n}(t) &:= \frac{U_{2,2,n}(t)}{0.98} + e_{2,n} \times \left(\frac{t^8}{8n^2} \left(\frac{K}{12} + \frac{P_{2,n}}{0.49}\right)^2 + \frac{0.621|t|^7 K^{3/4}}{12n^{3/2}} \left(\frac{K}{12} + \frac{P_{2,n}}{0.49}\right)\right), \\ P_{2,n} &:= \frac{1}{4} + \frac{P_{3,n}}{576}, \\ P_{3,n} &:= \frac{26.67 K^{1/2}}{n^{1/4}} + 4.8 \left(\frac{K}{n}\right)^{1/2} + \frac{1.24K}{n^{1/2}} + \frac{0.445K}{n^{3/4}} + \frac{0.04K}{n}, \\ e_{2,n} &:= \exp\left(0.01 \times \left(\frac{1}{6} + \frac{2P_{2,n}}{0.49}\right)\right), \\ U_{2,2,n}(t) &:= \frac{0.621|t|^5 K^{3/4}}{6n^{3/2}} + \frac{t^6 K}{24n^2} + \frac{t^6 0.621^2 K^{3/2}}{36n^2} + \frac{|t|^7 K^{7/4}}{72n^{5/2}} + \frac{t^8 K^2}{576n^3}. \end{split}$$

This results come from bounding from above the quantities $R_{2,n}(t,\varepsilon)$, $P_{2,n}(\varepsilon)$, $P_{3,n}(\varepsilon)$, $e_{2,n}(\varepsilon)$ and $U_{2,2,n}(t)$ that were defined in (Derumigny et al., 2021, p.32), with the choice $\varepsilon = 0.1$. We first upper bound $\widetilde{K}_{3,n}$ by $K_{3,n} + 1$ which itself is bounded by $K^{3/4} + 1$. For this, we use:

$$\begin{split} &-96\sqrt{2\varepsilon}\times 0.621<26.67,\\ &-32\times\varepsilon\times 0.621^2<1.235,\\ &-16\sqrt{2}\times 0.621\times\varepsilon^{3/2}<0.445,\\ &-(1-3\varepsilon)^2=0.49,\\ &-1.306\times 0.621^2<0.504,\\ &-e_3(\varepsilon)>1.0067 \end{split}$$

5.D.2 Definition of r_n^{cont}

We first show that the second term of Corollary 3.4 in Derumigny et al. (2021) is negligible and therefore can be included in the remainder term. Note that

$$0 \le \frac{1.0253 \kappa_n(\widetilde{K}_{3,n})^n \log(c_n)}{\pi} \le \frac{1.0253 \kappa_n(\widetilde{K}_{3,n})^n \log(8\pi^3 n^{3/2}/(t_1^* \widetilde{K}_{3,n}^3))}{\pi} \le \frac{1.0253 \kappa^n \log(8\pi^3 n^{3/2}/t_1^*)}{\pi} = O(n^{-5/4}),$$

since $c_n = 8\pi^3 n^{3/2} / (t_1^* \widetilde{K}_{3,n}^3)$ and $\widetilde{K}_{3,n} \ge 1$.

We now define r_n^{cont} as the latter term plus an upper bound on the expression that appears as $r_{2,n}$ in Corollary 3.4 in Derumigny et al. (2021):

$$\begin{aligned} r_n^{\text{cont}} &= \frac{1.0253 \,\kappa^n \log(8\pi^3 n^{3/2}/t_1^*)}{\pi} \\ &+ \frac{1.2533 (K^{3/4} + 1)^4}{16\pi^4 n^2} + \frac{0.208 K^{3/4} (K^{3/4} + 1)^4}{16\pi^4 n^{5/2}} \\ &+ \frac{14.1961 (K^{3/4} + 1)^{16}}{16^4 \pi^{16} n^8} + \frac{0.621 K^{3/4} \exp(-128\pi^6 n^4 / (K^{3/4} + 1)^8)}{3\pi \sqrt{n}} \\ &+ \frac{0.637 K^{3/4}}{3\pi \sqrt{2} \sqrt{n}} \left(\Gamma \left(3/2 \,, \, 0.1 \times (n/K)^{1/2} \right) - \Gamma \left(3/2 \,, \, T^{1/2} / 2\pi^2 \right) \right) + \\ &+ \frac{0.637 K^{3/4}}{3\pi \sqrt{2} \sqrt{n}} \left(\Gamma \left(3/2 \,, \, T^{1/2} / 2\pi^2 \right) - \Gamma \left(3/2 \,, \, T^{2} / 2\pi^2 \right) \right) \\ &+ J_3(T) + J_5(T) + \frac{K^{3/4}}{3\sqrt{n}} J_3(3, \sqrt{0.2} (n/K)^{1/4}, T^{1/4} / \pi, K^{3/4} + 1, T, n) \\ &+ \frac{0.504(e_{2,n} - 1.0067) K^{3/2}}{36n} + \frac{1.0253}{\pi} \int_0^{\sqrt{0.2} (n/K)^{1/4}} u e^{-u^2/2} R_{2,n}(u) du, \end{aligned}$$
(5.37)

where $\Gamma(\cdot, \cdot)$ stands for the incomplete Gamma function $\Gamma(a, x) := \int_x^{+\infty} u^{a-1} e^{-u} du$ and

$$J_5(T) := \frac{2}{T} \int_{T^{1/4}/\pi}^{T/\pi} |\Psi(u/T)| e^{-u^2/2} du,$$

with $\Psi(t) := \frac{1}{2} \left(1 - |t| + i \left[(1 - |t|) \cot(\pi t) + \frac{\operatorname{sign}(t)}{\pi} \right] \right) \mathbb{1} \{ |t| \le 1 \}.$

As in the definition of r_n^{gen} , we use the following inequalities to upper bound term-wise the expression $r_{2,n}$ of Derumigny et al. (2021):

- $|\lambda_{3,n}| \le 0.621 K_{3,n}$ ((Pinelis, 2011b, Theorem 1))
- $-K_{3,n} \leq K_{4,n}^{3/4}$ (convexity argument)
- $-\widetilde{K}_{3,n} \le K_{3,n} + 1;$
- $-0.3334 \times 0.621 < 0.208;$
- $-1.0253 \times 0.621 < 0.637$
- $1.306 \times 0.621^2 < 0.504,$
- $-e_3(\varepsilon) > 1.0067.$

We also use that the term $I_{5,2}(T)$ in Derumigny et al. (2021) is upper-bounded by

$$\frac{1.0253 |\lambda_{3,n}|}{3\pi\sqrt{2}\sqrt{n}} \Big(\Gamma\Big(3/2 \, , \, \varepsilon(n/K_{4,n})^{1/2} \Big) - \Gamma\Big(3/2 \, , \, T^{1/2}/2\pi^2 \Big) \Big),$$

setting here ε to 0.1 and using the aforementioned bound on $|\lambda_{3,n}|$.

Appendix 5.E Additional figures

As explained in the body of the article, this appendix presents the simulation results for the other data-generating process considered in Section 5.5.

The interpretation and main remarks are similar. Compared to the Gaussian case, the other DGP makes more visible the impact of replacing K_{ξ} by a plug-in instead of the fixed bound of 9: larger sample size are required to exit the \mathbb{R} regime and obtain informative $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n)$; when the interval is not shown in the graph, it is because we are in the regime $\operatorname{CI}_{u}^{\operatorname{Fin}}(1-\alpha,n) = \mathbb{R}$. In fact, for the Student scenario, the \mathbb{R} regime still prevails when n = 25,000 (the largest sample size considered in those simulation studies).

Figure 5.3: $K_{\xi} = 9 - \alpha = 0.05 - P^{\varepsilon | X}$ Laplace.



Figure 5.4: K_{ξ} by plug-in – $\alpha = 0.05 - P^{\varepsilon | X}$ Laplace.





Figure 5.5: $K_{\xi} = 9 - \alpha = 0.05 - \mathbf{P}^{\varepsilon \mid X}$ Gumbel.

Figure 5.6: K_{ξ} by plug-in $-\alpha = 0.05 - P^{\varepsilon | X}$ Gumbel.





Figure 5.7: $K_{\xi} = 9 - \alpha = 0.05 - \mathbf{P}^{\varepsilon \mid X}$ Student.



Titre : Contributions aux mesures de ségrégation ou de polarisation et à l'inférence non-asymptotique dans les modèles linéaires

Mots clés : inférence, non-asymptotique, modèles linéaires, ségrégation, polarisation, indices de ségrégation

Résumé : Cette thèse se compose de deux parties indépendantes. Les trois premiers chapitres s'intéressent aux mesures de ségrégation ou polarisation. Ces notions s'appliquent à des contextes variés mais partagent une formalisation commune. Une population est divisée en deux groupes exclusifs, un groupe dit minoritaire et un majoritaire, et les individus de ces deux groupes se répartissent entre des unités ou choisissent entre des options. Les unités sont par exemple des aires résidentielles, le groupe minoritaire les individus étrangers et le groupe majoritaire les individus français pour étudier un aspect de la ségrégation résidentielle en France. Cette modélisation permet également d'étudier la polarisation du langage politique en considérant comme unités ou options les entrées d'un dictionnaire et comme groupes minoritaire et majoritaire les parlementaires de différents partis ; ici, les individus choisissent les mots qu'ils utilisent. Qualitativement, il y a de la ségrégation ou polarisation si les deux groupes tendent à choisir systématiquement des options distinctes. En ségrégation résidentielle, les individus du groupe minoritaire sont concentrés dans certaines aires géographiques au lieu d'être répartis uniformément sur le territoire. Dans l'étude du langage politique, certains mots ou expressions seront davantage prononcés par un parti que par un autre. Quantitativement, la mesure de la magnitude de la ségrégation ou polarisation est confrontée au small-unit bias. Le premier chapitre présente un programme Stata qui implémente

trois méthodes permettant d'estimer des indices de ségrégation robustes au small-unit bias. Le second applique ces méthodes pour quantifier la ségrégation résidentielle en France entre 1968 et 2019 sur plusieurs dimensions (nationalité, statut social, position sur le marché du travail et proxy de l'ethnicité). Le troisième chapitre utilise les retranscriptions des débats au Congrès américain entre 1873 et 2016 pour étudier la polarisation du langage politique. La deuxième partie s'intéresse à la construction d'intervalles de confiance (IC) non-asymptotiques pour les coefficients des modèles de régression linéaire, un outil classique d'analyse économétrique. Ces intervalles reposent sur des majorations explicites de la distance entre la distribution empirique d'une somme normalisée de variables aléatoires indépendantes centrées admettant des moments d'ordre quatre et son expansion d'Edgeworth de 1er ordre. Ces majorations sont obtenues dans le quatrième chapitre, plus statistique et technique. Le cinquième et dernier chapitre utilise ces dernières pour construire des IC qui sont simultanément (i) valides pour toute taille d'échantillon (ii) sans imposer une distribution paramétrique de type Gaussienne ou l'indépendance entre régresseurs et résidus (et autorisent donc des résidus hétéroscédastiques), (iii) ayant une expression explicite, et (iv) atteignant asymptotiquement la même précision que les IC usuels fondés sur la normalité asymptotique de la statistique de Student.

Title: Contributions to measures of segregation or polarization and to nonasymptotic inference in linear models **Keywords:** inference, nonasymptotic, linear models, segregation, polarization, segregation indices

Abstract: This thesis is divided into two independent parts. The first three chapters deal with measures of segregation or polarization. The notion of segregation/polarization applies to various situations, but the formal modeling remains the same. Suppose that a population of individuals, comprised of a minority and a majority group, is allocated into units or makes choices over a set of options. When studying residential, occupational, or school segregation, units can be neighborhoods, firms, classrooms, and the minority group might be immigrants versus natives. The modeling also encompasses speech polarization, for instance. The US congresspeople are divided into Democrats and Republicans; the units or options are the items of a dictionary of words or phrases, and the speakers choose which words they use. Qualitatively, there is some segregation or polarization if the allocation or choice process leads to a situation where the two groups tend to select distinct units/options. In residential segregation, the minority individuals are concentrated in some units more than in others instead of being uniformly allocated. Regarding speech polarization, the presence of polarization means that Democrats and Republicans tend to use different words or phrases when they speak in Congress. Quantitative measures of segregation/polarization struggle with the socalled "small-unit bias." The first chapter presents a Stata command that implements three methods to estimate segregation indices robust to "the small-unit bias." The second applies these methods to

quantify residential segregation in France between 1968 and 2019 along various dimensions (nationality, social status, labor market position, proxy of races). Finally, the third chapter studies speech polarization in the US Congress between 1873 and 2016 using transcripts of congressional debates. The second part of this thesis is concerned with constructing nonasymptotic confidence intervals (CIs) for the individual coefficients of a linear regression model. Linear regression models are a ubiquitous method of econometric analyses. The CIs rely on explicit upper bounds on the uniform distance between the cumulative distribution function of a standardized sum of independent centered random variables with moments of order four and its first-order Edgeworth expansion. These bounds are derived in the fourth chapter, which is more technical and closer to statistics and probability than the other chapters. The last and fifth chapter uses these results to construct CIs that, at the same time, are (i) valid for any sample size (ii) without assuming parametric assumption such as the normality of error terms or independence between covariates and error terms (hence, our CIs allow for heteroskedasticity), (iii) have a closed-form expression, and (iv) whose length is asymptotically the same as the usual CI based on the t-statistic; thus our CIs have a coverage equals to the desired nominal level in the limit when the sample size goes to infinity.

